Chapter 3

Bounded Rationality

In the previous chapter, we appealed to a variety of notions of bounded rationality and simplicity to try to justify the special assumptions needed for the competitive limit theorem in Chapter 1. These included the Markov Property of equilibria used to characterize the equilibria of finite economies and the Continuity Principle imposed on the competitive sequences of equilibria.

If we take the idea of bounded rationality seriously, however, the complexity of these games is still very demanding. Furthermore, the assumption that the agents know not only their own equilibrium strategies but those of the other agents, is very demanding. Where do they get this information? Sometimes the common knowledge of equilibrium strategies is interpreted as the outcome of a process of introspective reasoning (Binmore (1990) calls this *eductive* reasoning). Sometimes it is treated as the outcome of a process of learning by trial and error. Clearly, the eductive approach does not reduce the computational ability required of the agents. The trial-and-error approach may do so. Adaptive, rule-of-thumb behavior is less demanding both informationally and computationally. If it leads to equilibrium behavior, it may provide some support for the notion that boundedly rational individuals can acquire strategies that are close to equilibrium strategies.

In this chapter, I present an example of this kind of rule-of-thumb or adaptive behavior that leads not very bright agents to a competitive equilibrium. There are many examples of rule-of-thumb or adaptive behavior in the economics literature and it is not my intention to review them all here; a few relevant papers are mentioned in the next section. One of the problems with this kind of model of boundedly rational behavior is that there are so many possibilities and so few accepted modeling principles. This is why the charge of "ad hocery" is so often and so justifiably aimed at this kind of theorizing. There is no a priori defense against these charges: the proof of the pudding is in the eating. If this attempt provides some insight or some surprises, then it may not have been a waste of time.

3.1 Imitation and Experimentation

The model that is the focus of this chapter is based on ideas that were first explored in a paper I wrote with Robert Rosenthal. Before describing the model, I begin with a brief summary of the model and results from Gale and Rosenthal (1998), hereafter referred to as GR. GR studies the "learning" behavior of boundedly rational agents who play a strategic game repeatedly over time. The broad objective of this line of research is to see whether boundedly rational agents can learn to play the equilibrium strategies of the game. The answer to this question turns out to be complex and subtle. Much of the interest of the paper lies in the complex dynamics that are generated by apparently simple behavioral rules.

The GR model can be seen as an application of bounded rationality to social learning. The process of learning to play this game is "social" in two senses. First, because of strategic interaction through the game, one agent's learning (adaptation) affects the learning of the others. Secondly, agents can learn from each other by imitating the behavior that they observe in the rest of the population.

The social aspect of learning raises a number of well known efficiency issues. One of these is the *free-rider problem* associated with informational externalities (Chamley and Gale (1994) and Caplin and Leahy (1994)). In models of asymmetric information, an agent's actions reveal his private information. Because the agent is only concerned with his own payoff, he ignores the value of this information to other agents. This externality typically leads to inefficient decisions: in equilibrium, either too little information is revealed or it is revealed too slowly.

Another aspect of social learning that has attracted a lot of attention is the phenomenon of *herd behavior*. In some models of herd behavior, (e.g., Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)), agents ignore their own information and base their decisions on the public information revealed by the actions of their predecessors. This information may be incorrect, in which case the decisions based on it will be inefficient. More importantly, the agents who decide to join the herd are suppressing their own information. Because their actions are independent of their own information, they do not reveal their private information. The result may be that only a small fraction of the information available to the agents as a group ever becomes common knowledge. Even in models where agents never ignore their own information completely, it can be shown that under certain circumstances there is always a positive probability that agents herd on the incorrect choice (Smith and Sorensen (1996)). In models with endogenous timing (Chamley and Gale (1994), and Gul and Lundholm (1992)), the inefficiency of social learning can take the form of delay rather than herd behavior.

The literature on social learning and herd behavior is firmly in the tradition of rational, Bayesian, maximizing behavior (with some exceptions, e.g., Ellison and Fudenberg (1993, 1995)). In a model of bounded rationality, there is no explicit decision to be a free rider or to join a herd. The behavior rules of the agents are assigned exogenously. However, free riding and herd behavior do have an anologue in the imitative behavior that is crucial to the dynamics of the GR model.

The GR model focuses on two types of behavior, imitation and experimentation. Imititators copy what they see other agents doing. Experimenters try new strategies randomly and persist with the strategies that do best. While a sensible agent might engage in both experimentation and imitation, GR simplifies by assuming that each agent specializes in exactly one of these activities. In a further simplification, GR assumes that there is only one experimenter and that the rest of the agents are imitators.

There is a large literature dealing with optimal experimentation. Some, such as Banks and Sundaram (1992) and Bolton and Harris (1998), are models of rational learning, rather than rule-of-thumb learning, and they assume that the underlying environment is stationary. Others, such as Aghion, Bolton and Harris (1991), are closer in spirit to the random experimenter in the GR model. The motivation given by GR for assuming random search is that agents have very little information about the environment in which they are operating and have very limited ability to process the information they do have. In that situation, the least demanding strategy is to search randomly. The fact that search is random turns out to have important implications for the dynamics of the model, as we shall see.

Another important feature of the model is that the search for a better strategy goes on indefinitely. In most models of learning, there is a single fixed parameter that agents try to estimate. As time passes, the agents' beliefs converge to the true value of this parameter, their estimates become insensitive to new information, and experimentation dies out. This is true of the literature on learning in games, for example, including models of fictitious play and Bayesian learning (Fudenberg and Kreps (1993), Jordan (1993), Kalai and Lehrer (1993a,b), Krishna and Sjostrom (1995), Marimon (1995), and Benaim and Hirsch (1996)). The GR model, by contrast, assumes that experimentation continues indefinitely.

The interest in the case of permanent experimentation comes from the observation that we live in a non-stationary environment. When the environment is constantly changing, one can never assume that one is close to the equilibrium or that experimentation can stop. Models of fictitious play, Bayesian learning, and adaptive learning assume both a stationary environment and that agents place less and less weight on recent experience as time passes. Such increasing inertia is essential to guarantee convergence. It makes sense in a stationary environment, where individual behavior can in principle converge to an equilibrium and beliefs can converge to the truth. In a world that is constantly changing, agents have no reason to assume that they have reached a permanent state of equilibrium. Consequently, they do have reason to continue to experiment and to give significant weight to recent experience. Although GR studies a stationary environment in order to obtain clean and transparent results, the model is motivated by the assumption that agents always have something to learn.

In each period, the imitators observe their own actions and the actions of the other agents. Then they adjust their actions a constant fraction λ of the distance between their previous actions and the average action of the other agents. Since it is only the average of the imitators' actions that matters and their decision rule is linear, there is no loss of generality in replacing them with a representative agent.

At each date $t = 1, 2, ...$ the agents play a symmetric normal-form game. They are all asummed to have the same payoff function, but only the experimenter makes use of the payoff to update his action. If q is the action chosen by the experimenter and $\bar{q}$ is the average action of the imitators, then the experimenter's payoff is

$$-(q - B\overline{q})^2,$$

B is the slope of each player's best-response line. When $B < 0$, agents' actions are strategic substitutes, as in the standard Cournot model; when

$B > 0$, actions are strategic complements. When $B \neq 1$, the game has a unique symmetric equilibrium, in which all agents choose $q = 0$.

Clearly, if the imitators choose $\bar{q}$, the best response is $B\bar{q}$. Because of the quadratic form of the payoff function, a strategy is better for the experimenter if and only if it is closer to the best response $B\bar{q}$ than his current strategy q. The experimenter searches randomly for a better strategy. If the experimenter chose a strategy q_{t-1} last period, then he "tests" a new strategy that is uniformly distributed on the interval $[q_{t-1} - 1, q_{t-1} + 1]$. Note that this interval is centered on last period's action and that its size does not change over time. If the randomly drawn strategy falls in the better-response set then he adopts it as q_t. Otherwise, he puts $q_t = q_{t-1}$.

Denote by X_t the experimenter's action at t and by Y_t the average of the imitators' actions at t. Then, under the behavioral rules specified above, for any initial condition (x_0, y_0) these behavioral rules define a Markov chain $\{(X_t, Y_t)\}$ having state space $\mathbf{R}^2$.

This model has a number of interesting dynamic properties.

- First, assuming $B < 1$, it is *stable in the large*. This means that from any initial state (x_0, y_0) the chain converges with probability one to a compact neighborhood of the origin.

- Secondly, for the case of strategic substitutes (B negative and sufficiently large in absolute value) the symmetric equilibrium is *unstable in the small*. This means that for any sufficiently small neighborhood of the origin and any initial condition (x_0, y_0) in that neighborhood ($x_0 \neq 0$), the chain leaves the neighborhood with probability one.

- Finally, for the case of strategic substitutes (B negative and sufficiently large in absolute value), it is *not too unstable*. This means that, for any neighborhood of the equilibrium, however small, the probability of the chain being in the given neighborhood at date t converges to one as t approaches ∞.

At the macroscopic level, stability in the large tells us that these adaptive rules do work, by bringing the agents' actions to a compact neighborhood of the equilibrium where they are approximately optimal. On the other hand, at the microscopic level, instability in the small tells us that the agents can never learn the equilibrium strategies exactly. Their actions fluctuate permanently around the equilibrium levels. How large these fluctuations are depends on the parameters λ and B and the size of the search window (here normalized to two).

The third result is puzzling, since it appears to contradict the second. The two results can be reconciled, but the explanation depends on another feature of the model, namely, that the chain $\{(X_t, Y_t)\}$ is *null recurrent*. Although the chain leaves sufficiently small neighborhoods of the origin with probability one, the expected time this takes is infinite. The reason is that for states (x, y) very close to the origin the better-response set is very small. It is hard to find a better response and, as a result, the chain changes very slowly. If an econometrician looked at the cross-sectional frequency distribution of a set of sample paths, he might conclude that the process was converging. This is a result of the fact that almost every path occasionally comes close to the origin and then takes a very long time to get away. If the same observer looked at a single sample path, he might come to a very different conclusion.

Null recurrence is a strange property: it seems to lie between convergence and instability. It depends on the fact that close to the equilibrium nothing much happens most of the time, although when something does happen the result is "unstable". To test the robustness of the null recurrence result, GR studies a version of the game with small random perturbations to the payoffs. The payoff function of the experimenter in the perturbed game is

$$-(q_t - B\overline{q}_t - \varepsilon_t)^2$$

for every t, where $\{\varepsilon_t\}$ is an i.i.d. sequence that takes the values ε and $-\varepsilon$ with probability $1/2$ each. The shock ε_t simply shifts the best response function up or down by a small amount. For small values of ε the behavior of the resulting chain is similar to that of the unperturbed chain. Stability in the large and instability in the small continue to hold in the perturbed model, but the not-too-unstable result changes. The long-run behavior of the chain is described by a non-degenerate invariant probability measure and the probability distribution of states at date t converges to the invariant distribution as $t \to \infty$.

Two properties of the model are crucial for these results: the presence of strategic substitutes and the non-vanishing size of the search window. To get a better sense of the importance of strategic substitutes, GR also studies the (unperturbed) model under the assumption that the game exhibits strategic complements. More precisely, they assume that $0 < B < 1$, so the strategic complements are "not too strong." Under this assumption the symmetric equilibrium is shown to be stable in a strong sense: for any initial condition (x_0, y_0) the chain converges to the origin with probability one.

To examine the role of the non-vanishing window size, GR extend the base-case model to allow for exogenously shrinking search windows. Formally, they assume that the experimenter chooses a strategy randomly from the interval $[x_{t-1} - d_t, x_{t-1} + d_t]$, where $d_t > 0$, for every t. Here again they find that the equilibrium is stable in the same strong sense when the size of the search window converges to 0 as long as it does not converge too fast. More precisely, if

$$\lim_{t\to\infty} d_t = 0 \text{ and } \lim_{T\to\infty} \sum_{t=1}^{\infty} d_t = \infty$$

then for any initial condition (x_0, y_0) the chain converges to 0 with probability one. This result holds whenever $B < 1$.

What have we learned from these exercises?

- The first lesson is that the interaction of two simple types of behavioral adaptation can produce endogenous cycles. This may turn out to be a useful way of looking at certain kinds of macroeconomic fluctuations.

- A second lesson concerns the roles of strategic complements and substitutes. There has been a lot of interest in using models with strategic complements to explain the severity of macroeconomic fluctuations. In such models if individual activity levels are strategic complements, each agent's best response is an increasing function of the activity levels of the others. If an agent increases his activity because of an exogenous shock, the others will increase their actions too. In this way, strategic complementarity magnifies the effect of the inital shock to the economy. One of the interesting features of our model is that strategic substitutes are necessary for local instability, whereas with strategic complements the model is very stable.

- A third lesson is that the cycles in the models with strategic substitutes have a highly structured complexity that does not appear in other models in the literature. For example, in simulations we find that the amplitude of these cycles varies over time, sometimes being very damped and then growing again; but these variations are regular in the sense that the average amplitude of successive cycles is positively correlated.

- A fourth lesson concerns the role of experimentation. The randomness in the experimenter's behavior is essential to generate the changing relative rates of adaptation (between experimenters and imitators) that drive the dynamics. When the chain is close to the equilibrium, the experimenter rarely finds a better strategy than the one that he is currently using. When the chain is far from the equilibrium, he finds a better strategy relatively frequently. Thus, adaptation is faster or slower for the experimenter depending on the degree of disequilibrium in the system. The imitators, on the other hand, are constantly adjusting their actions toward the average action (hence, on average, toward the experimenter's action) at a constant proportionate rate. This means that the relative speeds of adjustment for experimenter and imitators vary, depending on the distance from equilibrium; and that explains how the model can be stable in the large but unstable in the small. It also accounts for the fact that the system can spend long periods close to the equilibrium, then cycle away in an increasing orbit for a long period, and then approach the equilibrium again.

In the rest of this chapter, I present a simple model of boundedly rational behavior in the context of a market for a single good. The essential idea is to assume that agents search at random for a better strategy, where a strategy is a limit price at which the agent is willing to trade the good. The motivation for random search is the same as in the GR paper: in the first place, if agents knew exactly where to look for a better strategy, they would not be learning at all; secondly, boundedly rational agents cannot master the computational complexity of Bayesian learning. So, we are left with the assumption that they search randomly for better strategies. In this version of the model, there is no role for imitative behavior, though imitative behavior can easily be introduced.

3.2 A Behavioral Model of Competition

3.2.1 The Market

There is a single indivisible commodity, called "the good", that can be exchanged in integer amounts and there is a divisible numeraire commodity, called "money". Traders are divided into buyers and sellers. There are N

sellers indexed by $i = 1, ..., N$ and N buyers indexed by $j = 1, ..., N$. There is no loss of generality in assuming equal numbers of buyers and sellers, since an agent with an extreme valuation is effectively not a participant in the market.

Preferences are quasi-linear and each agent wants to buy or sell at most one unit of the good. This means that each agents' preferences can be parameterized in terms of their valuation of one unit of the good. The valuation of seller i is denoted by u_i and the valuation of buyer j is denoted by v_j. If a seller i has x_i units of the good and m_i units of money, his utility is

$$U_i(x_i, m_i) = u_i x_i + m_i.$$

Similarly, the utility of a buyer with y_j of the good and m_j units of money is

$$U_j(y_j, m_j) = v_j y_j + m_j.$$

In practice, the quasi-linearity of the utility function allows us to normalize initial holdings of the good and money to zero and henceforth conduct the analysis in terms of surplus or "gains from trade". (For analytical simplicity we ignore the non-negativity constraints on quantities of the good and money). If seller i exchanges one unit of the good for p units of money, his surplus is $p - u_i$. Similarly, if buyer j obtains one unit of the good for p units of money, his gain from trade is $v_j - p$. In this notation, we can put

$$U_i(-x_i, p) = (p - u_i)x_i$$

and

$$U_j(y_j, -p) = (v_j - p)y_j,$$

where $x_i, y_j \in \{0, 1\}$.

The primitive data of the market, then, are the size of the market, N, and the valuations $u = (u_1, ..., u_N)$ and $v = (v_1, ..., v_N)$ of the sellers and buyers, respectively. The agents are ordered so that $0 < u_1 < u_2 < ... < u_N$ and $v_1 > v_2 > ... > v_N > 0$. We also assume that no buyer and no seller have the same valuation:

$$u_i \neq v_j, \forall i, j = 1, ..., N.$$

The assumption that no two agents have the same valuation simply rules out inconvenient ties. The assumption is satisfied generically in the sense that if the market data (N, u, v) were chosen at random, the probability of violating these assumptions would be zero.

3.2.2 Market-Clearing Prices

For any market data, (N, u, v), we have to distinguish four possible configurations of valuations. The *marginal seller* (resp. *marginal buyer*) is the highest valuation seller (resp. lowest valuation buyer) who gets to trade in a competitive equilibrium. The index of the marginal agent, which is the same for sellers and for buyers, of course, is denoted by m and defined by the conditions that

$$u_m < v_m \text{ and } u_{m+1} > v_{m+1}.$$

The sellers $i = 1, ..., m$ (resp. buyers $j = 1, ..., m$) who trade in a competitive equilibrium are called *infra-marginal.*

Because of the discreteness of demand and supply (each agent wants to trade zero or one units), there exists a non-degenerate interval of market-clearing prices that equate demand and supply. For any market data (N, u, v), exactly one of the following four configurations is (generically) possible:

$$\begin{array}{ll} (A-B) & u_m < v_{m+1} < u_{m+1} < v_m \\ (A'-B') & v_{m+1} < u_m < v_m < u_{m+1} \\ (A'-B) & v_{m+1} < u_m < u_{m+1} < v_m \\ (A-B') & u_m < v_{m+1} < v_m < u_{m+1}. \end{array}$$

In each case, there is a different set of market-clearing prices. Only the first m agents on each side of the market can trade so the price must be less than or equal to u_{m+1} to exclude seller $m+1$ and greater than or equal to v_{m+1} to exclude buyer $m+1$. In Case $(A-B)$, for any price in this interval, precisely the first m agents on each side of the market want to trade, so the set of market-clearing prices is the interval $[v_{m+1}, u_{m+1}]$. By similar reasoning, we can calculate the interval of market-clearing prices in each case:

$$\begin{array}{ll} (A-B) & [v_{m+1}, u_{m+1}] \\ (A'-B') & [u_m, v_m] \\ (A'-B) & [u_m, u_{m+1}] \\ (A-B') & [v_{m+1}, v_m]. \end{array}$$

The interval of market-clearing prices can be more compactly denoted by $[c_0, c_1]$, where $c_0 = \max\{u_m, v_{m+1}\}$ and $c_1 = \min\{u_{m+1}, v_m\}]$. A situation in which all infra-marginal agents trade at a single price belonging to this interval will be referred to as a *perfectly competitive outcome.*

3.2.3 The Market Game

The trading process is represented by a normal-form stage game in which agents submit limit orders to a profit-maximizing market maker, who then arranges trades between pairs of buyers and sellers.

The agents' strategies are the limit prices at which they are willing to trade one unit. Each seller i chooses an *asking price* a_i and each buyer j chooses a *bid price* b_j. The asking price a_i signifies that the seller is willing to supply a single unit of the good at any price equal to a_i or higher. Similarly, the bid price b_j signifies that the buyer is willing to purchase one unit of the good at any price up to and including b_j.

For simplicity, the agents' strategies are restricted to a compact interval. Seller i's asking price a_i is restricted to be at least as great as his private valuation u_i and no greater than some large finite number M. By restricting the sellers' strategies in this way, we are assuming that no agent is so unintelligent as to choose a dominated strategy. This seems reasonable and is not particularly important in what follows. Buyer j's bid price b_j is restricted to be non-negative and no greater than his valuation v_j. Let

$$\mathbf{X} = \{(a, b) \in \mathbf{R}_+^N \times \mathbf{R}_+^N | \, u_i \leq a_i \leq M, 0 \leq b_j \leq v_j, \forall i, j\}$$

denote the set of strategy profiles for the buyers and sellers and denote a typical strategy profile by $x = (a, b)$ where $a = (a_1, ..., a_N)$ and $b = (b_1, ..., b_N)$ are the strategy profiles for sellers and buyers, respectively.

Suppose that the agents submit the limit orders (a, b). The market-maker arranges matching trades to maximize his profits. By submitting limit orders, the agents have committed themselves to trading at any price that does not violate their limits. In order to maximize his profit, the market maker will execute the trades at the limit prices, that is, he will pay seller i the tendered ask price a_i and charge buyer j the tendered bid price b_j. Denote seller i's trade by ξ_i, where $\xi_i = 0$ means that i does not trade and $\xi_i = 1$ means that i sells one unit, and denote buyer j's trade by ζ_j, where $\zeta_j = 0$ means that j does not trade and $\zeta_j = 1$ means that j buys one unit. Then the market maker's profit will be

$$\Pi(\xi, \zeta, a, b) = \left\{ \sum_{j=1}^{N} b_j \zeta_j - \sum_{i=1}^{N} a_i \xi_i \right\},$$

where $\xi = (\xi_1, ..., \xi_N)$ and $\zeta = (\zeta_1, ..., \zeta_N)$. Formally, the market maker's problem is to choose $(\xi,\zeta) \in \{0,1\}^{2N}$ to solve

$$\Pi^*(a,b) = \max_{(\xi,\zeta)} \Pi(\xi,\zeta,a,b)$$

subject to the usual feasibility constraint

$$\sum_{j=1}^{N} \zeta_j \leq \sum_{i=1}^{N} \xi_i.$$

A generic profile $(a,b) \in \mathbf{X}$ is one in which no two agents, whether both buyers or both sellers or one buyer and one seller, choose the same limit price. For any generic profile (a,b), there is a unique profit-maximizing assignment of trades (ξ,ζ) defined by the conditions that $\sum_{i=1}^{N} \xi_i = \sum_{j=1}^{N} \zeta_j$, $\xi_{i_0} > \xi_{i_1}$ implies that $a_{i_0} < a_{i_1}$, $\zeta_{j_0} > \zeta_{j_1}$ implies that $b_{j_0} > b_{j_1}$, and $\xi_i = \zeta_j = 1$ implies that $a_i < b_j$. Hence, for any generic profile (a,b) in $\mathbf{X}$, this trading mechanism defines a unique payoff $\pi_i^s(a,b)$ for seller i and $\pi_j^b(a,b)$ for buyer j. (When the context makes it clear whether the trader is a buyer or seller, the superscripts s and b are dispensed with).

To ensure a unique outcome (ξ,ζ) for each strategy profile (a,b), I assume that where the market-maker is indifferent between two offers, bids or asks, he randomizes between the agents with equal probabilities. Where he is indifferent between executing a trade and not executing it, I assume that he executes it. These assumptions simply serve to define a unique trading mechanism and are not crucial in what follows.

In what follows, there is no loss of generality in restricting the discussion to generic profiles. The reason is that buyers and sellers choose their strategies randomly, so the probability of observing a non-generic profile is zero. This completes the definition of the normal-form game $\Gamma = (N, \mathbf{X}, \pi)$.

Several aspects of this matching and trading procedure are noteworthy:

- The market-maker executes the trades at the bid and ask prices, keeping the difference as his profit.
- The procedure maximizes the volume of trade, subject to the constraint that each trade be voluntary.
- The procedure is efficient in the sense that trades go to the buyers who offer the most and the sellers who demand the least.

The fact that there is a maximizing agent at the center of the market is significant. For example, if we had adopted an alternative approach, matching pairs of buyers and sellers in each period and letting them bargain over the terms of trade, none of the three features above would necessarily hold. The use of a market-maker thus introduces by design an element of efficiency that is not characteristic of all trading mechanisms.

The second and third properties are implications of the market maker's maximizing behavior. Since the market maker can choose the price at which to execute a trade, subject to the limits of the traders' orders, he will want to arrange as many trades as possible as long as he makes a non-negative profit on each one. Similarly, profit maximization implies that the buyers with the highest bids and the sellers with the lowest asks will get to trade. A more decentralized procedure, such as random matching and bargaining, would not necessarily have these very useful properties.

3.2.4 Behavioral Rules

In this section, I describe the behavior of the individual traders. Rather than assigning beliefs to agents and assuming that each agent maximizes his long-run payoff relative to these beliefs, I define behavioral rules directly for the agents. The motivation for this approach is the realization that maximization is too demanding in many contexts. It is interesting to see whether "simpler" behavioral rules can lead agents to use equilibrium strategies. Of course, the environment we are studying is itself very simple, a reflection of our own bounded rationality. So in order to capture the notion that agents have limited ability to understand complex systems—that is, agents are simple relative to their environment—we shall have to assume that they are very stupid indeed.

There are many ways of specifying boundedly rational behavior—as was pointed out in Section 3.1, this is one of the weaknesses of the approach—and the behavioral rules specified here are not the only ones that might recommend themselves for study. But they are simple and provide a vehicle for discussing a number of interesting issues. The basic idea is that agents of very limited intelligence seach at random for good strategies and abandon their current strategies when they encounter a better one. Random search presumes no knowledge of the environment (the structure of the model) apart from the appropriate strategy set. It requires no memory apart from the knowledge of the agent's current strategy. And it requires no foresight, in fact, it assumes complete myopia on the part of the agents. This is about the simplest kind of behavior one could imagine and for that reason it is a natural place to start.

Before getting down to the details of the search procedure, one aspect deserves some discussion. The idea of random search employed here involves choice (between the current strategy and a randomly selected alternative) and that implies that the agent "knows" the payoff from both strategies. How does the agent learn the payoffs? The eductive approach assumes that the agent knows the payoff functions, so that he can calculate the payoffs from any strategy; but then there is nothing to stop him from calculating the best response rather than searching at random. I do not want to assume that the agent knows the payoff function because it assumes a degree of sophistication and "computing power" that is incompatible with random search.

A less demanding interpretation is that the agent learns the payoff by experimenting with different strategies. Imagine that every so often a single agent gets a chance to search for a better strategy. This involves selecting a strategy at random, trying it for a short period of time, noting the payoff flow, comparing that flow to the flow from the previous strategy and then choosing between the new and the previous strategy based on the payoff comparison. While all this is going on, the other agents continue to use the same strategy. This procedure does not demand too much "brain power" on the part of the agent, but it is rather cumbersome to describe and makes the dynamics rather complicated. So instead I assume that agents can conduct "virtual experiments", in which they discover the payoff to a new strategy instantaneously and adopt it only if the payoff is greater than or equal to that of their current strategies. This is an approximation to a real experiment with a new strategy, which lasts for a short but finite interval. Nothing of importance seems to hang on this simplification.

Another point to note about the search procedure is that only one agent at a time is allowed to search. If two or more agents were experimenting at the same time, the result of their experiments with new strategies might be misleading. For example, a strategy that appeared to be better for seller i when buyer j was experimenting with a new strategy might actually turn out to be worse when buyer j decides to return to his previous strategy. These kinds of errors are ruled out by assuming that only one agent at a time can experiment with a new strategy. If search is a costly and hence discrete event, the probability of two agents searching at the same time is likely to be small. Here we take the probability to the limit and assume it is zero. This does not mean that search is slow, since the time scale is arbitrary. However, it does mean that there is an element of inertia in the system. A related assumption of inertia is found in the evolutionary game literature, where it is sometimes assumed that only a fraction of a population is allowed to change strategies in any period (e.g., Kandori, Mailath and Rob (1993) and Young (1993)).

Trade takes place at a sequence of dates $t = 1, 2, \ldots$. At each date, one of the $2N$ agents is chosen at random to alter his strategy. If seller i is chosen, he randomly chooses a new strategy from his strategy set X_i. If this price gives a (weakly) higher payoff than his current ask price, he adopts the new price as his strategy. If not, he retains his previous strategy. Similarly, if a buyer j is chosen he randomly chooses a new strategy from his strategy set X_j. If this price gives a (weakly) higher payoff, he adopts the new bid price as his strategy. Otherwise, he retains the existing strategy.

These simple rules define a stochastic process. The state vector at any date is the current profile of limit order strategies. Suppose that $x = (a, b)$ is the state at date $t - 1$ and $x' = (a', b')$ is the state at the next date t. If seller i is chosen to move at date t, a new ask price ω_i is chosen according to the uniform distribution on $[u_i, M]$. Then $b' = b$ and

$$a' = \begin{cases} (\omega_i, a_{-i}) & \text{if } \pi_i((\omega_i, a_{-i}), b) \geq \pi_i(a, b) \\ a & \text{if } \pi_i((\omega_i, a_{-i}), b) < \pi_i(a, b). \end{cases}$$

Similarly, if buyer j is chosen to move at date t, then ω_j is drawn from a uniform distribution on $[0, v_j]$ and the new state satisfies $a' = a$ and

$$b' = \begin{cases} (\omega_j, b_{-j}) & \text{if } \pi_j(a, (\omega_j, b_{-j})) \geq \pi_j(a, b) \\ b & \text{if } \pi_j(a, (\omega_j, b_{-i})) < \pi_j(a, b). \end{cases}$$

Note that in the event that agents are indifferent between the new strategy and the existing strategy, they switch to the new. In particular, if the initial state $x = (a, b)$ is one in which seller i cannot trade, so that $\pi_i(a, b) = 0$, then any price in X_i is weakly preferable and so the new price is drawn randomly from X_i. This assumption is important because it eliminates the possibility of getting stuck in situations where no trade is possible. For example, suppose that $a_i = M$ for every seller i and $b_j = 0$ for every buyer j. This is a possible state and one in which no trade is possible. Moreover, no deviation by any single buyer or seller will make trade possible. So if agents only switch to *strictly* better strategies, the original position will be a rest point for the system and there is no possibility of convergence to a competitive outcome.

For this reason, it is essential to allow agents to switch to *weakly* better strategies.

3.2.5 The Markov Chain $\{X_t\}$

[The material in this subsection is rather technical and the reader who is not interested in this detail may wish to skip ahead to Section 3.3, noting only Lemma 1 and its corollary, which are used in the sequel.]

A stochastic process is a family $\{X_t\}$ of random elements, defined on an underlying probability space $(\Omega, \mathcal{F}, P)$. A *Markov chain* $\{X_t\}$ is a particular type of stochastic process, with a countable parameter set $T = \{1, 2, ...\}$, a sequence of random elements X_t taking values in the state space $\mathbf{X}$, and a probability distribution P_x that satisfies:

$$P_x(X_t \in A | X_{t_1} = x_{t_1}, ..., X_{t_k} = x_{t_k}) = P_x(X_t \in A | X_{t_k} = x_{t_k})$$

for any initial state x and for any times $t_1 < ... < t_k < t$.

The behavioral rules described in the preceding section define a Markov chain. In fact, the evolution of the stochastic process can be represented by a (time-invariant) transition probability $P(x, A)$ which tells us, for any current state x, the probability that the state of the system at the next date belongs to the set A. The purpose of this section is to translate the behavioral rules defined above into a formal definition of the transition probabilty $P(x, A)$.

The set of admissible strategy profiles $\mathbf{X}$ is endowed with the Borel σ-field $\mathcal{B}(\mathbf{X})$ of measurable sets (the σ-field generated by the open sets of $\mathbf{X}$). Using the rules outlined above, for any given initial position, $x = (a, b)$ and any measurable set $A \in \mathcal{B}(\mathbf{X})$, we can define a Markov transition probability $P(x, A)$. For any agent k and any state x, let $B_k(x)$ denote the *better than set*, that is,

$$B_k(x) = \{x'_k | \pi_k(x'_k, x_{-k}) \geq \pi_k(x)\}.$$

Since the agent searches randomly for a better strategy, if he finds a better strategy it will be uniformly distributed on the better than set. Then let $U_k(x, \cdot)$ denote a probability distribution on $\mathbf{X}$ such that the support of $U_k(x, \cdot)$ is $[B_k(x) \times \{x_{-k}\}]$ and the restriction of $U_k(x, \cdot)$ to the set $[B_k(x) \times \{x_{-k}\}]$ is the uniform distribution. For practical purposes we can think of $U_k(x, \cdot)$ as the uniform distribution on $[B_k(x) \times \{x_{-k}\}]$ but it is important to remember that it is defined for every measurable subset of $\mathbf{X}$. Let $D(x, \cdot)$ denote the Dirac distribution concentrated on x and let $\delta_k(x)$ denote the ratio of the diameter of the better response set to the strategy set of agent k:

$$\delta_k(x) = \begin{cases} \mathrm{diam} B_k(x)/(M - u_k) & \text{if } k \text{ is a seller} \\ \mathrm{diam} B_k(x)/v_k & \text{if } k \text{ is a buyer.} \end{cases}$$

Then, conditional on the current state x and agent k being chosen to move, the distribution of the new state is given by

$$G_k(x, \cdot) = \delta_k(x) U_k(x, \cdot) + (1 - \delta_k(x)) D(x, \cdot).$$

With probability $\delta_k(x)$ agent k finds a better strategy and, given that he finds a better strategy, his strategy is uniformly distributed on $B_k(x)$; with probability $1 - \delta_k(x)$ he does not find a better strategy and, if he does not find a better strategy, the new state will be the same as the old. Since each of the agents is chosen to move with equal probablity, the transition probability can be defined by putting

$$P(x, A) = (2N)^{-1} \sum_{i=1}^{N} G_i(x, A) + (2N)^{-1} \sum_{j=1}^{N} G_j(x, A),$$

for each state x and measurable set A. In other words, we take $G_k(x, A)$ to be the transition probability conditional on k being chosen to search for a better strategy; and then the (unconditional) transition probability $P(x, A)$ is just the expected value of the conditional transition probabilities $\{G_k(x, A)\}$, where the weights N^{-1} are the probabilities of choosing each agent as the experimenter.

If $P : \mathbf{X} \times \mathcal{B}(\mathbf{X}) \to \mathbf{R}_+$ is the transition probability function for a Markov chain then by definition it satisfies the following two conditions:
(i) for each $A \in \mathcal{B}(\mathbf{X})$, $P(\cdot, A)$ is a non-negative, measurable function on $\mathbf{X}$;
(ii) for each $x \in \mathbf{X}$, $P(x, \cdot)$ is a probability measure on $\mathcal{B}(\mathbf{X})$.
From the definition of the function P it is clear that property (ii) is satisfied. To show that property (i) is satisfied, we only need to note that for any measurable subset $A \subset \mathbf{X}$, the probability $G_k(A)$ is a measurable function of x.

For any initial state x, the transition probability $P(x, A)$ defines a Markov chain $\{X_t\}$ on the probability space $(\Omega, \mathcal{F}, P_x)$. In this case, we take the set Ω to be $\mathbf{X}^\infty$, the countable product of copies of $\mathbf{X}$, $\mathcal{F}$ to be the σ-field generated by the cylinder sets $A_1 \times ... \times A_t \in \mathcal{B}(\mathbf{X}) \times ... \times \mathcal{B}(\mathbf{X})$ and P_x to be the unique extension of the set functions defined inductively on the cylinder sets $A_1 \times ... \times A_t$ by the transition kernel P:

$$\begin{aligned} P_x^1(x, A_1) &= P(x, A_1) \\ P^2(A_1 \times A_2) &= \int_{A_1} P(x, dy_1) P(y_1, A_2) \\ &\vdots \\ P^n(A_1 \times ... \times A_t) &= \int_{A_1} P(x, dy_1) \int_{A_2} P(y_1, dy_2) ... P(y_{t-1}, A_t). \end{aligned}$$

[See Meyn and Tweedie (1993) , Section 3.4, for the details of the procedure and a proof.] This Markov chain has the property that for any initial condition x and any cylinder set $A_1 \times ... \times A_t$

$$P_x(X_1 \in A_1, ..., X_t \in A_t) = \int_{A_1} P(x, dy_1) \int_{A_2} P(y_1, dy_2) ... P(y_{t-1}, A_t).$$

A stochastic process with finite dimensional distributions satisfying this property for every t is called a *time homogenous Markov chain* with transition probability kernel $P(x, A)$ and initial condition x.

There are a couple of technical results that will be used repeatedly in the sequel and which it will be convenient to state here.

For any set $A \in \mathcal{B}(\mathbf{X})$ and initial condition x, let $L(x, A)$ denote the probability that the Markov chain $X = \{X_t\}$ enters the set A at some date t, that is,

$$L(x, A) = P_x[X \in A],$$

where $[X \in A]$ denotes the event $\{\omega | X_t(\omega) \in A, \exists t\}$. Let $Q(x, A)$ denote the probability that it enters A infinitely often, that is,

$$Q(x, A) = P_x[X \in A, \text{ i.o}],$$

where $[X \in A, \text{i.o}]$ denotes the event $\{\omega | X_t(\omega) \in A, \text{ for infinitely many } t\}$. We say that a set $B \in \mathcal{B}(\mathbf{X})$ is *accessible* from a set $A \in \mathcal{B}(\mathbf{X})$ if $L(x, B) > 0$ for every initial condition $x \in A$ and we say that B is *uniformly accessible* from A if there exists a $\delta > 0$ such that

$$\inf_{x \in A} L(x, B) \geq \delta.$$

If B is uniformly accessible from A we write $A \rightsquigarrow B$. A set $A \in \mathcal{B}(\mathbf{X})$ is called *Harris recurrent* if $Q(x, A) = 1$ for any $x \in A$. The next result tells us that if A is uniformly accessible, then the chain visits A infinitely often with probability one. Intuitively, if the probability of going from A^c to A is at least δ then the probability of staying out of A forever must be $(1-\delta)^\infty = 0$. Then second part tells us that if B is uniformly accessible from A then if A is visited infinitely often, so is B.

Lemma 1 *(i) For any $A \in \mathcal{B}(\mathbf{X})$, if $\mathbf{X} \rightsquigarrow A$ then A is Harris recurrent; in fact, $Q(x, A) = 1$ for any $x \in \mathbf{X}$; (ii) For any sets $A, B \in \mathcal{B}(\mathbf{X})$, if $A \rightsquigarrow B$ then $\{X \in A \text{ i.o.}\} \subset \{X \in B \text{ i.o.}\}$ a.s.*

[Meyn and Tweedie (1993), Theorem 9.1.3.] A set $A \in \mathcal{B}(\mathbf{X})$ is called *absorbing* if $P(x, A) = 1$ for all $x \in A$. In other words, once in A the system stays there forever with probability one. The next result tells us that if B is absorbing and B is uniformly accessible from a disjoint set A, then the probability of visiting A infinitely often is zero; because once the system enters B it never returns and if it visits A infinitely often it must visit B infinitely often, contradicting the Lemma.

Corollary 2 *For any sets $A, B \in \mathcal{B}(\mathbf{X})$ satisfying $A \cap B = \emptyset$, if B is absorbing and $A \rightsquigarrow B$ then $Q(x, A) = 0$ for any x.*

Proof. Let E_t be the event $\{X_t \in B,\ X_s \in A, \exists s > t\}$. Since B is absorbing, $P_x(E_t) = 0$ for all t and since the probability measure P_x is σ-additive, $P_x(\cup_{t=1}^{\infty} E_t) = 0$. From conclusion (ii) of the preceding Lemma,$\{X \in A \text{ i.o.}\} \subset \{X \in B \text{ i.o.}\}$, so

$$\{X \in A \text{ i.o.}\} \subset \cup_{t=1}^{\infty} E_t.$$

Hence, $Q(x, A) = P_x\{X \in A \text{ i.o.}\} = 0$ as required.

We are now ready to study the convergence properties of the Markov chain $\{X_t\}$.

3.3 Convergence to Competitive Prices

3.3.1 Volume-maximizing trade

The main objective is to show that, in some sense, the Markov chain described above converges to a competitive outcome, that is, to a situation in which all trade takes place at competitive prices. The analysis of convergence can be conveniently broken down into two steps. The first step involves showing that, with probability one, the volume of trade is maximized in finite time and, from that point onwards, the volume of trade remains constant. The next step is to show that, asymptotically, all trade takes place at market-clearing prices.

To analyze the volume of trade, we first need some additional notation. For any state $x \in \mathbf{X}$ and any measurable set $A \in \mathcal{B}(\mathbf{X})$, define $P^n(x, A)$ to be the probability that the system belongs to A exactly n periods later, given that it started at x, that is, $P^n(x, A) = P_x(X_n \in A)$. Then $P^n(x, A)$ can be defined recursively by putting

$$P^n(x, A) = \int P(x, dy)P^{n-1}(y, A).$$

Let A be the set of (generic) states such that m units of the good are traded whenever the state of the system belongs to A. Let A' denote the subset of A in which it is the infra-marginal agents $i = 1, ..., m$ and $k = 1, ..., m$ who get to trade. Then we show that A' is accessible from any state in $\mathbf{X}$. In mathematical notation this is written $\mathbf{X} \rightsquigarrow A'$. Although the proof is rather lengthy, accessibility simply requires one to show that, from any starting point, there is a positive probability (bounded away from zero) that random search will lead the traders to a configuration of strategies (prices) that is consistent with maximal trade.

Lemma 3 $\mathbf{X} \rightsquigarrow A'$.

Proof. In fact, we can show that for any initial condition x there exists an integer $n > 0$ and a number $\alpha > 0$ such that $P^n(x, A') \geq \alpha$.

First, let $x = (a, b)$ and take the infra-marginal sellers $i = 1, ..., m$ and arrange them in decreasing order according to their asking prices. That is, let $\{i_1, ..., i_m\}$ be an ordered m-tuple of the first m sellers such that $1 \leq i_h \leq m$ for any h and

$$a_{i_h} > a_{i_{h+1}}, \forall h = 1, ..., m - 1.$$

There is a probability $(2N)^{-m}$ that the sellers are chosen to move in precisely this sequence in the periods $t = 1, ..., m$.

Suppose that the sellers are chosen in just this sequence to choose new strategies. When seller i_h has a chance to move, he draws a new price according to the uniform distribution from the interval $[u_{i_h}, M]$ and with probability

$(\bar{a} - c_0)/(M - u_{i_h})$ the new asking price a'_{i_h} lies in the interval $[c_0, \bar{a})$, where $\bar{a} = (c_0 + c_1)/2$. The probability of this sequence of events is at least

$$(2N)^{-m} \left(\frac{(\bar{a} - c_0)}{(M - u_1)} \right)^m$$

since seller 1 has the smallest probability of picking a price in the interval. According to our behavioral rules, if seller i_h weakly prefers the new price he has chosen, he adopts it. Otherwise, he sticks with the old price.

Notice that once we have gone through the sequence of infra-marginal sellers and determined a new N-tuple of asking prices $a' = (a'_1, ..., a'_m, a_{m+1}, ..., a_N)$, it must be the case that for each $i = 1, ..., m$ either seller i can trade at the new prices (a', b) or $a'_i \in [c_0, \bar{a})$. To see this, note that $a'_i \notin [c_0, \bar{a})$ only if a_i is weakly preferred to the prices in $[c_0, \bar{a})$, that is, i can already trade when he has the opportunity to change his strategy and $a_i > \bar{a}$. Since asking prices are decreasing in this sequence, it follows that all the sellers in $\{i_{h+1}, ..., i_m\}$ can trade when their turn to move comes. Since these sellers can already trade, a change in their prices will not affect the ability to trade of seller i.

Now go through an exactly similar argument with the buyers. Arrange the m infra-marginal buyers in order of increasing bid prices. Then the set $\{j_1, ..., j_m\}$ satisfies $1 \leq j_h \leq m$ for all h and

$$b_{j_h} < b_{j_{h+1}}, \forall h = 1, ..., m - 1.$$

With probability $(2N)^{-m}$ the m buyers are chosen in precisely this order and, when it is buyer j_h's turn to move, with probability $(\bar{a}, c_1]/v_{j_h}$ he chooses a new price in the interval $(\bar{a}, c_1]$. Of course, he picks the new price as his strategy if and only if it is weakly preferred to b_{j_h}, given the choices of the other agents. The probability of this sequence of events is at least

$$(2N)^{-m} \left(\frac{(c_1 - \bar{a})}{v_1} \right)^m$$

since buyer 1 has the smallest probability of choosing a price in the interval $(\bar{a}, c_1]$.

Once we have allowed all the infra-marginal buyers to choose a new strategy, we note that for every buyer $j = 1, ..., m$ either $a'_j \in (\bar{a}, c_1]$ or buyer j can trade at the state $x' = (a', b')$. Again, this follows from the fact that if $b'_j \in (\bar{a}, c_1]$ then $b_j < \bar{a}$ and buyer j was able to trade at b_j when he had the opportunity to move. Since all the other buyers were offering higher prices, a change in their prices cannot effect buyer j.

I claim now that all infra-marginal buyers and sellers must trade at (a', b'). The sellers who were able to trade at the original prices can still do so, since changes in the buyers' prices will not affect the sellers' ability to trade. Likewise, buyers who were able to sell at their original prices can still do so. The rest of the infra-marginal sellers $i = 1, ..., m$ (resp. buyers $j = 1, ..., m$) are charging prices $a'_i \in [c_0, \bar{a})$ (resp. $b'_j \in (\bar{a}, c_1])$. It is clear that all the extra-marginal sellers $i = m+1, ..., N$ (resp. buyers $j = m+1, ..., N$) must charge prices $a_i < c_0$ (resp. $b_j > c_1$), so the market maker will maximize profits by arranging trades among the rest of the infra-marginal buyers and sellers.

Thus, after $n = 2N$ periods, with probability at least

$$\alpha = (2N)^{-2m} \left(\frac{(\bar{a} - c_0)}{(M - u_1)} \right)^m \left(\frac{(c_1 - \bar{a})}{v_1} \right)^m,$$

$X_n \in A'$. ■

Since $A' \subset A$, Lemma 3 implies that A is accessible from $\mathbf{X}$. Then Lemma 1 implies that from any initial condition x, the system almost certainly reaches a state in which the volume of trade is maximal, in finite time.

Lemma 4 *For any initial state x, X reaches A in finite time with probability one, that is $L(x, A) = 1$ for any $x \in \mathbf{X}$.*

In fact, Lemma 1 tells us something stronger, namely that $Q(x, A) = 1$ for any x, but we do not need this result since we can show that the set A is absorbing: once X reaches A it remains there with probability one.

Lemma 5 *The set A is absorbing.*

Proof. Suppose that $x = (a, b) \in A$ is a generic initial state and let x' denote the immediately following state along some realization. If the agent who is chosen to move at date 1 is an agent who is already trading, then the same set of agents are trading in the new state x'. This follows because the moving agent will never change his price unless it is (weakly) better for him to do so and this requires that he trade. The identities of the other trading agents are unaffected, because the moving agent is merely moving his price within the trading range.

If the agent who is chosen to move is a non-trading agent, then either his new choice does not fall in the trading range, in which case the set of trading agents is unchanged, or it does fall into the trading range, in which case he replaces one of the trading agents on his side of the market. However, the total number of trading agents cannot be changed. Since we are already at the maximum number of trades, the number of trades could only fall and this will not happen because all the existing trades continue to be profitable. ■

3.3.2 Convergence to a competitive outcome

Now that we have established the inevitability of maximal trade, we can show that the chain converges almost surely to a competitive outcome. The analysis can be restricted to the set A of generic states in which m units of the good are traded. Again, however, this requires us to examine a number of different cases. We characterize these cases in terms of the *bid-ask spread.* First, at any state $x \in A$, define the *marginal ask* $\alpha(x)$ and the *marginal bid* $\beta(x)$ by putting

$$\alpha(x) = a_{i_m} \text{ and } \beta(x) = b_{j_m},$$

where the sellers $i_1, ..., i_m, ..., i_N$ are ordered by increasing asking prices and the buyers $j_1, ..., j_m, ..., j_N$ are ordered by decreasing bid prices. The marginal ask is the highest price demanded by a seller that is accepted in state x; similarly, the marginal bid is the lowest price offered by a buyer that is accepted in state x. These are just m-th order statistics for the ask and bid prices, arranged in increasing and decreasing order, respectively. Clearly, $\alpha(x) < \beta(x)$ for any (generic) state $x \in A$.

For any initial state x in A, put

$$(\alpha_t, \beta_t) \equiv (\alpha(X_t), \beta(X_t)), \forall t = 1, 2, ...$$

The sequences $\{\alpha_t\}_{t=1}^{\infty}$ and $\{\beta_t\}_{t=1}^{\infty}$ are stochastic processes that have some useful properties. As noted above, $\alpha_t > \beta_t$ for all t with probability one. Furthermore,

$$\alpha_t < u_{m+1} \Longrightarrow \alpha_t \leq \alpha_{t+1}$$

and

$$\beta_t > v_{m+1} \Longrightarrow \beta_t \geq \beta_{t+1}.$$

To see this, consider the position of the sellers when $\alpha_t < u_{m+1}$. Since there are m trading sellers and α_t is the highest accepted asking price, the sellers who manage to trade must be $i = 1, ..., m$. At date $t+1$ either none of these sellers changes his price, in which case $\alpha_{t+1} = \alpha_t$, or one of them is chosen to move and finds a better price, which must be a higher one. This price change either leaves the marginal ask unchanged or raises it, so in either case $\alpha_{t+1} \geq \alpha_t$. The explanation of the second implication is similar. These monotonicity properties are used to prove the following convergence theorem. Let B denote the set of states such that both marginal prices are contained in the competitive interval $[c_0, c_1]$, that is,

$$B = \{x \in A |\, c_0 < \alpha(x) < \beta(x) < c_1\}.$$

Theorem 6 *Suppose that the initial state $x \in B$. Then there exists a random variable X_∞ such that*

$$\lim_{t\to\infty} \alpha_t = \lim_{t\to\infty} \beta_t = X_\infty \in (c_0, c_1), P_x\text{-a.s.}$$

and, for any seller $i = 1, ..., m$, $X_{it} \to X_\infty$ almost surely (resp. for any buyer $j = 1, ..., m$, $X_{jt} \to X_\infty$ almost surely).

Proof. By the monotonicity properties of $\{\alpha_t\}$ and $\{\beta_t\}$, $\lim_{t\to\infty} \alpha_t$ and $\lim_{t\to\infty} \beta_t$ both exist and are measurable and $\lim_{t\to\infty} \alpha_t \leq \lim_{t\to\infty} \beta_t$. Furthermore, we can show that for any $\varepsilon > 0$ there is a $\delta > 0$ such that $|\alpha_t - \beta_t| > \varepsilon$ implies that with probability at least δ, $|\alpha_{t+1} - \beta_{t+1}| < \varepsilon/2$. To see this, note that with probability $m(2N)^{-1}$ one of the sellers $i = 1, ..., m$ is chosen to move and with probability $\varepsilon/2(M - u_i)$ he chooses a price in the interval (α_t, β_t). Then it is clear that the probability that $|\alpha_t - \beta_t| > \varepsilon$ for all t is zero, for any fixed but arbitrary $\varepsilon > 0$. Hence, $\lim_{t\to\infty} \alpha_t = \lim_{t\to\infty} \beta_t = X_\infty \in (c_0, c_1)$ as required.

To show that the prices of the individual buyers and sellers also converge to c, one uses similar arguments. For example, if $|X_{it} - \alpha_t| > \varepsilon$ for all t, for any fixed $\varepsilon > 0$, then we can find a fixed $\delta > 0$ such that with probability δ, $|X_{it+1} - \alpha_{t+1}| < \varepsilon/2$. Hence the probability that $|X_{it} - \alpha_t| > \varepsilon$ for all t and any fixed $\varepsilon > 0$ must be zero and this, together with the monotonicity of $\{X_{it}\}$ implies that $X_{it} \to X_\infty$ almost surely. The proof for buyers is exactly similar. ■

This convergence theorem gives us the essential argument for establishing the competitive outcome, for in the limit we have all infra-marginal agents charging the same (random) price X_∞ which belongs to the competitive interval $[c_0, c_1]$. However, the assumption that the initial state is such that both marginal prices belong to the interval $[c_0, c_1]$ leaves something to be done. How do we know that such a situation will eventually arise? In fact, in some cases, the marginal bid and ask prices may only approach the competitive interval asymptotically. For this reason the competitive convergence theorem takes the following form.

Theorem 7 *For any initial condition x, with probability one, either (i) X reaches B in finite time, in which case Theorem 6 tells us that the chain converges to a competitive outcome, or (ii) X does not reach B in finite time, but α_t and β_t converge to a common limit belonging to the set $\{c_0, c_1\}$.*

Establishing part (ii) turns out to be rather tedious and involves ruling out a number of other possible cases. In each case, the key fact is again the monotonicity of $\{\alpha_t\}$ and $\{\beta_t\}$.

Case 1.
The first case to consider is the possibility that the interval $[\alpha_t, \beta_t]$ contains the competitive interval in its interior infinitely often. Let C denote the set of states such that $[c_0, c_1] \subset [\alpha(x), \beta(x)]$, that is,

$$C = \{x \in A | \, \alpha(x) < c_0 < c_1 < \beta(x)\}.$$

The next result shows that $\{X_t\}$ belongs to C infinitely often with probability 0. The idea of the lemma is to show that B is an absorbing state and that B is uniformly accessible from C.

Lemma 8 *For any initial condition x, $Q(x, C) = 0$.*

Proof. Theorem 6 implies that B is an absorbing set, so it is sufficient to note that $C \rightsquigarrow B$ and then apply Lemma 2. To see that B is uniformly accessible from C, note that with probability $(2N)^{-1}$ the seller $i = 1$ is chosen at date $t+1$ and with probability $(c_0 + c_1)/2(M - u_1)$ he chooses a price in the interval $(c_0, (c_0 + c_1)/2)$; similarly, with probability $(2N)^{-1}$ the buyer $j = 1$ is chosen at date $t+1$ and with probability $(c_0 + c_1)/2v_1$ he chooses a price in the interval $((c_0 + c_1)/2, c_1)$. ■

Case 2.

Using the result from Case 1, we can show that α_t (resp. β_t) must be greater than c_0 (resp. less than c_1) infinitely often. If not then one of the marginal prices α_t and β_t lies outside the competitive interval $[c_0, c_1]$ for all sufficiently large values of t and then Lemma 8 implies that the other price must lies inside.

Lemma 9 *For any initial condition x, with probability one, one of the following must occur:*

(i) $c_0 < \alpha_t < \beta_t < c_1$ for some t;
(ii) $\alpha_t < c_0$ for all t sufficiently large;
(iii) $\beta_t > c_1$ for all t sufficiently large.

Proof. Suppose that case (i) does not occur. Then since $\alpha_t < \beta_t$ for all t, it must be the case that $\alpha_t < c_0$ infinitely often or $\beta_t > c_1$ infinitely often. Since the marginal prices do not straddle the competitive interval infinitely often, except for a finite number of dates $\alpha_t < c_0$ implies that $\beta_t < c_1$ and $\beta_t > c_1$ implies that $\alpha_t > c_0$. Then having ruled out case (i), $\alpha_t < c_0$ (resp. $\beta_t > c_1$) implies that $\alpha_{t+1} < c_0$ (resp. $\beta_{t+1} > c_1$) except for a finite number of dates. Thus, (ii) and (iii) are the only alternatives to (i). ■

Case 3.
If case (i) of Lemma 9 holds, then we are done. It remains to show that cases (ii) and (iii) do not rule out the competitive outcome. If one of these cases does occur, the marginal price converges to the appropriate endpoint, c_0 or c_1, without ever entering the interval $[c_0, c_1]$. Then the prices at which trade occurs will be competitive in the limit, but the marginal trade can oscillate indefinitely between the marginal trader and an extra-marginal trader. For example, along some realization (sample path) of the Markov chain, we may have $\alpha_t < c_0$ for all t sufficiently large; then $\alpha_t \to c_0$ as $t \to \infty$. Unfortunately, we cannot analyze individual sample paths because they occur with probability zero. Instead we have to reformulate the claim in terms of a set of states.

For any $\varepsilon > 0$ let A_ε denote the set of states $x = (a, b)$ such that

$$a_i \leq \min\{v_m, u_{m+1}\} + \varepsilon \text{ for } i = 1, ..., m$$

and

$$b_j \geq \max\{u_m, v_{m+1}\} - \varepsilon \text{ for } j = 1, ..., m.$$

The next lemma says that, starting from any initial condition x and for any ε, A_ε is uniformly accessible from any point in the state space $\mathbf{X}$. Unfortunately, while the proof of the lemma is straightforward, it is rather tedious. It consists of considering arbitrary initial states and showing inductively that with positive probability, price-cutting by sellers and price-raising by buyers will lead the infra-marginal traders to choose prices close to the competitive interval.

Lemma 10 *For any* $\varepsilon > 0$, $\mathbf{X} \rightsquigarrow A_\varepsilon$.

Proof. More precisely, we shall show that for any x and $\varepsilon > 0$, there exists an integer $n > 0$ (independent of ε and x) and a number $\gamma > 0$ such that $P^n(x, A_\varepsilon) \geq \gamma$. We begin with the sellers and then repeat the argument for the buyers. We need to distinguish two different cases, depending on whether $v_m > u_{m+1}$ or $u_{m+1} > v_m$.

Case A. $v_m > u_{m+1}$. In this case, let $x = (a, b)$ be the initial condition and consider the first $m + 1$ sellers, $i = 1, ..., m + 1$. Arrange them in descending order according to their asking prices: $i = i_1, ..., i_{m+1}$ so that $a_{i_h} > a_{i_{h+1}}$ for $h = 1, ..., m$. Note that we only consider generic initial states, in which there are no ties. Since ties occur with probability zero, this does not entail any loss of generality in what follows. Now suppose that these agents are chosen to make changes in their strategies in precisely this order. This will happen with probability $(2N)^{-(m+1)} > 0$. Consider first seller i_1. If $a_{i_1} \leq u_{m+1} + \varepsilon$ there is nothing more to prove so suppose that $a_{i_1} > u_{m+1} + \varepsilon$. At this price, the agent is unable to trade, since at most m sellers can trade, the market maker always assigns trades to the sellers with the lowest asking prices, and there are m sellers with lower prices. Therefore this seller, when given the chance, will randomize on the interval $[u_{i_1}, M]$ and has a positive probability $(u_{m+1} + \varepsilon - u_{i_1})/(M - u_{i_1}) > 0$ of choosing a price $a_{i_1} \leq u_{m+1} + \varepsilon$.

Now move to the second seller. If this seller is choosing a price $a_{i_2} \leq u_{m+1} + \varepsilon$, there is nothing more to prove. If $a_{i_2} > u_{m+1} + \varepsilon$ then seller i_2 has the highest price and cannot trade. Then he will randomize on the interval $[u_{i_2}, M]$ and there is a probability $(u_{m+1} + \varepsilon - u_{i_2})/(M - u_{i_2}) > 0$ of choosing a price satisfying $a'_{i_2} \leq u_{m+1} + \varepsilon$.

To complete the proof by induction, suppose that $a'_i \leq u_{m+1} + \varepsilon$ for $i = i_1, ..., i_h$ and consider i_{h+1}. If $a_{i_{h+1}} \leq u_{m+1} + \varepsilon$ there is nothing to prove so suppose that $a_{i_{h+1}} > u_{m+1} + \varepsilon$. Then i_{h+1} cannot trade at this price, (since there are m lower prices), so when given the chance seller i_{h+1} will randomize on the interval $[u_{j_{h+1}}, M]$ and with probability $(v_{m+1}+\varepsilon-u_{i_{h+1}})/u_{i_{h+1}}$ the new price will satisfy $a'_{i_{h+1}} \le u_{m+1}+\varepsilon$. By induction we have shown that for all $i = i_1, ..., i_{m+1}$ there is a positive probability that $a_i \le u_{m+1}+\varepsilon$ within exactly $m+1$ periods. In fact, the probability of this happening is at least

$$\frac{(2N)^{-(m+1)}\varepsilon^{m+1}}{(M-u_{m+1})^{m+1}}$$

since seller $m+1$ has the smallest probability of choosing a price in the acceptable interval.

Case A$'$. $u_{m+1} > v_m$. In this case, arrange the first m sellers $i = 1, ..., m$ in decreasing order according to their asking prices: $i_1, ..., i_m$ satisfy $a_{i_h} > a_{i_{h+1}}$ for $h = 1, ..., m-1$. Then there is a probability $(2N)^{-m}$ that these sellers will be chosen in precisely this order to change their asking prices. If $a_{i_1} > u_{m+1}$ he cannot trade since there are $m-1$ sellers with lower prices and there are only $m-1$ buyers who are willing to buy at prices as high as a_{i_1}. Consequently, this seller will randomize on the interval $[u_{i_1}, M]$ and choose a price satisfying $a'_{i_2} \le v_m$ with probability $(v_m - u_{i_1})/(M - u_{1_1})$. Suppose that all sellers $i = i_1, ..., i_h$ have chosen a price $a'_i \le v_m$ and that $a_{i_{h+1}} > v_m$. Then i_{h+1} cannnot trade at his current price. When he gets a chance to change his strategy he will randomize over the interval $[u_{i_{h+1}}, M]$ and choose a price $a_{i_{h+1}} \le v_m$ with probability $(v_m - u_{i_{h+1}})/u_{i_{h+1}}$. By induction, we have shown that all sellers $i = 1, ..., m$ will choose a price less than or equal to v_m in exactly m periods with a probability greater than or equal to

$$\frac{(2N)^{-m}(v_m-u_m)^m}{(M-u_{m+1})^m}.$$

We proceed in an exactly similar way with the buyers, this time distinguishing the cases $u_m > v_{m+1}$ and $u_m < v_{m+1}$.

Case B. $v_{m+1} > u_m$. In this case, buyer $m+1$ can sometimes compete with buyer m for trade. Take the first $m+1$ buyers $j = 1, ..., m+1$ and arrange them in increasing order of bid prices $j_1, ..., j_{m+1}$ so that $b_{j_h} < b_{j_{h+1}}$ for all $h = 1, ..., m$. There is a probability $(2N)^{-(m+1)}$ that the buyers will be chosen in precisely this order to change their strategies. Consider the decision of buyer j_1. Since he is offering the lowest price of $m+1$ buyers, he cannot possibly trade. Therefore his new strategy will be a random draw from the interval $[0, v_{j_1}]$. With probability $(v_{j_1} - v_{m+1}+\varepsilon)/v_{j_1}$ the new price will satisfy $b'_{j_1} \ge v_{m+1}-\varepsilon$. Suppose that $b'_j \ge v_{m+1}-\varepsilon$ for $j = j_1, ..., j_h$

and consider j_{h+1}. If $b_{j_{h+1}} \geq v_{m+1} - \varepsilon$ there is nothing to prove so suppose that $b_{j_{h+1}} < v_{m+1} - \varepsilon$. Then j_{h+1} cannot trade at this price, (since there are m higher prices), so when given the chance buyer j_{h+1} will randomize on the interval $[0, v_{j_{h+1}}]$ and with probability $(v_{j_{h+1}} - v_{m+1} + \varepsilon)/v_{j_{h+1}}$ the new price will satisfy $b'_{j_{h+1}} \geq v_{m+1} - \varepsilon$. By induction we have shown that for all $j = j_1, ..., j_{m+1}$ there is a positive probability that $b_j \geq v_{m+1} - \varepsilon$ within $m+1$ periods. In fact, the probability of this event will be at least

$$\frac{(2N)^{-m+1}\varepsilon^{m+1}}{{v_{m+1}}^{m+1}},$$

since buyer $j = m+1$ has the smallest probability of choosing an acceptable price.

Case B′. $v_{m+1} < u_m$. Arrange the first m buyers in increasing order according to their bid prices: $j_1, ..., j_m$ satisfy $b_{j_h} < b_{j_{h+1}}$ for $h = 1, ..., m-1$. Then there is a probability $(2N)^{-m}$ that these sellers will be chosen in precisely this order to change their asking prices. If $b_{j_1} < u_m$ he cannot trade since there are $m-1$ sellers with lower prices and there are only $m-1$ sellers who are willing to accept a price as low as b_{j_1}. Consequently, this seller will randomize on the interval $[0, u_{j_1}]$ and choose a price satisfying $b'_{j_1} \geq u_m$ with probability $(u_m - v_{j_1})/v_{j_1}$. Suppose that all buyers $j = j_1, ..., j_h$ have chosen a price $b'_j \geq u_m$. Then j_{h+1} cannnot trade at his current price. When he gets a chance to change his strategy he will randomize over the interval $[0, v_{j_{h+1}}]$ and choose a price $b_{j_{h+1}} \geq u_m$ with probability $(u_m - v_{j_{h+1}})/v_{j_{h+1}}$. By induction, we have shown that all buyers $j = 1, ..., m$ will choose a price $b'_j \geq u_m$ in m periods with probability greater than or equal to

$$\frac{(2N)^{-m}(u_m - v_m)^m}{{v_m}^m}.$$

The concatenation of these two arguments, for sellers and buyers respectively, obviously establishes the desired result. Note however that in Cases A and B, for sellers or buyers, the probability of reaching the set A_ε depends on the value of ε. This follows from the fact that we use the extra-marginal seller or buyer $m+1$ to discipline the infra-marginal sellers and buyers, respectively, and his chance of choosing an acceptable price is proportional to ε. ■

Because Lemma 10 holds for arbitrary $\varepsilon > 0$, we conclude that even in cases (ii) and (iii) the limit points of *one* of the marginal prices, $\{\alpha_t\}$ or $\{\beta_t\}$, are contained in the competitive interval $[c_0, c_1]$.

Lemma 11 *Suppose that $\alpha_t < c_0$ (resp. $\beta_t > c_1$) for all t sufficiently large; then $\lim_{t\to\infty} \alpha_t = c_0$ (resp. $\lim_{t\to\infty} \beta_t = c_1$).*

Proof. Let A_n be the set A_ε with $\varepsilon = 1/n$ and let

$$\Omega_n = \{\omega \in \Omega | \, X_t \in A_n \text{ i.o.}\},$$

where "i.o." means for all but a finite number of t. Since $P_x(\Omega \backslash \Omega_n) = 0$,

$$\begin{aligned} P_x(\cap_{n=1}^{\infty} \Omega_n) &= P_x(\Omega \backslash \cup_{n-1}^{\infty} (\Omega \backslash \Omega_n)) \\ &\geq 1 - P_x(\cup_{n-1}^{\infty} (\Omega \backslash \Omega_n)) \\ &\geq 1 - \sum_{n=1}^{\infty} P_x(\Omega \backslash \Omega_n) = 1. \end{aligned}$$

Let $\Omega' = \cap_{n=1}^{\infty} \Omega_n$ and let

$$\Omega'_0 = \{\omega \in \Omega' | \, \alpha_t < c_0 \text{ i.o.}\}$$

and

$$\Omega'_1 = \{\omega \in \Omega' | \, \beta_t > c_1 \text{ i.o.}\}.$$

Then by construction we know that for any $\omega \in \Omega'_0$ (resp. any $\omega \in \Omega'_1$), $\alpha_t(\omega) \nearrow c_0$ (resp. $\beta_t(\omega) \searrow c_1$), as required. ■

By a similar argument, we can show that in both cases (ii) and (iii), *both* marginal prices converge.

Lemma 12 *Suppose that $\alpha_t < c_0$ (resp. $\beta_t > c_1$) for all t sufficiently large; then $\lim_{t\to\infty} \beta_t = c_0$ (resp. $\lim_{t\to\infty} \alpha_t = c_1$) where as usual the limits hold almost surely.*

Proof. For any $\varepsilon > 0$ let B_ε denote the set of states such that $\alpha(x) < c_1 - \varepsilon$ and $\beta(x) > c_1$, that is,

$$B_\varepsilon = \{x \in A | \, c_0 < \alpha(x) < c_1 - \varepsilon < c_1 < \beta(x)\},$$

where A is the set of states consistent with maximum trade. For any state $x \in B_\varepsilon$ there is a positive probability, bounded away from zero, of entering the set $B = \{x \in A | \, c_0 < \alpha(x) < \beta(x) < c_1\}$. For example, with probability $m/2N$, an infra-marginal buyer is chosen to move and with probability at least ε/v_1 chooses a price in the interval $[c_1 - \varepsilon, c_1]$. Once in the set B, the system never leaves, so we can use a previous argument to show that the system cannot be in the set B_ε infinitely often: for any initial state x and any fixed by arbitrary $\varepsilon > 0$,

$$P_x(X_t \in B_\varepsilon \text{ i.o.}) = 0.$$

Since we have also shown already that the system cannot be in a straddling position infinitely often, it follows that, defining $B'_\varepsilon = \{x \in A | \, \alpha(x) < c_1 - \varepsilon < c_1 < \beta(x)\}$, we have

$$P_x(X_t \in B'_\varepsilon \text{ i.o.}) = 0.$$

From this it is immediate that, on a subset of Ω having full measure, if case (iii) occurs then $\alpha_t \to c_1$.

By an exactly similar argument, show that in case (ii) $\beta_t \to c_0$. ■

So, in both cases (ii) and (iii), we have shown that the marginal prices converge to elements of the competitive interval $[c_0, c_1]$. Because the marginal prices only approach the competitive interval asymptotically, it is possible that the extra-marginal agents $i = m + 1$ and $j + m + 1$ are able to trade infinitely often. However, this trade will be infrequent because the probability that an extra-marginal agent is able to trade in any period goes to zero as the marginal prices converge to c_0 or c_1.

3.4 Extensions

3.4.1 Finite Strategy Sets

Although we have been analyzing a simple model, the analysis turns out to be complicated. Some of the complexity of the analysis in this chapter is associated with the fact that the state space X is an infinite set. It is natural to ask if the analysis could be simplified by replacing X with a finite subset.

For example, if the agents were restricted to choosing from a finite number of bid and ask prices, the set of admissible strategy profiles would be finite. The behavior of the system could then be represented by a finite Markov chain and this would eliminate some of the technical complications.

There is a certain realism in the assumption that admissible prices must belong to a finite set. In practice, the price of an indivisible goods will be an integral multiple of the smallest monetary unit (cents, lire, etc.) and since the price is naturally bounded in this context, this implies that only a finite number of prices needs to be considered.

Assuming a finite number of prices simplifies the analysis in some ways and makes it more complicated in others. With a continuum of prices, the probability of two agents quoting the same price is zero, so we do not have to worry about ties. If prices have to be an integral number of monetary units, then it is natural to assume that each agent can choose from the same set of prices and this makes ties a probable event. The treatment of ties may be important for the evolution of the system. For example, if goods are allocated randomly in the event of a tie, it may make it difficult for the system to settle down unless one agent can undercut the other. Suppose that two buyers have very similar reservation values and that there is no price that lies between the two values. Then both could end up offering the same price and the trade would bounce between them until one of them raised his bid.

One way to avoid ties is to assume that different agents have different sets of admissible prices. This seems a little contrived, given the motivation for finiteness in terms of a minimal currency unit, but it would work. The assumption that no two agents have the same reservation value plays a similar role in the preceding model.

There is another problem that arises even if the set of admissible prices is fine enough so that there exists an admissible price between any two reservation values. An example will make this clear. Suppose that there is a single buyer with a reservation price of 3 and two sellers, one with a reservation price of 1 and the other with a reservation price of 2. Call these agents $1, 2$ and 3. Then a competitive equilibrium price must lie in the closed interval $[1, 2]$. Suppose that the smallest admissible price greater than 2 is 2.1. If 2.1 belongs to the price set of seller 1, then seller 1 can hold the marginal price and there is nothing seller 2 can do about it. On the other hand, if 2.1 belongs to the price set of seller 2, then at some point seller 1 will have to choose an ask price less than two in order to trade and in that case the marginal price will eventually get inside the competitive interval. Whether convergence occurs or not depends on exactly how we choose the admissible prices.

If the grid of admissible prices were sufficiently fine, one might hope the outcome would be approximately competitive. In other words, we can prove a limit theorem characterizing the asymptotic prices as $t \to \infty$ and then prove another limit theorem characterizing these prices as the grid became finer and finer until it approximated the continuous state space X. The fact that, for some choices of the grid, the asymptotic prices (as $t \to \infty$) remain outside the competitive interval echoes case (ii) of Theorem 7.

Taking limits provides us with an anologue of Theorem 7 but is unlikely to be any easier and in some ways could be more complicated than dealing directly with the continuous state space X.

3.4.2 Trading Multiple Units

One of the important simplifications of the theory developed in this chapter is the assumption that each agent wants to trade a single unit of the good. There are markets where this is a reasonable simplification (think of the labor market or the market for consumer durables, for example) and there are many examples in the literature on economic theory where substantial progress has been made with the aid of this assumption. Nonetheless, it is a restrictive assumption and it limits the applicability of the theory.

There is one particular feature of the assumption that should worry us given, our interest in perfect competition. As long as an agent only wants to trade at most one unit of the good, there is absolutely no incentive for the agent to distort his demand or supply in order to influence the price. The agent only has two alternatives, to trade or not to trade. If he does not trade his payoff is zero and the price is irrelevant. So an agent must trade to get a positive payoff. He will try to trade at the best possible price but there is no useable tradeoff between the amount traded and the price.

The fact that we get a competitive outcome for any finite number of agents in the market is surprising. One might suspect that the assumption that each agent trades only one unit has something to do with it. For this reason it is important to examine the case of a market in which agents can trade multiple units of the good.

When an agent wants to trade more than one unit, he immediately faces the possibility of a tradeoff between price and quantity. This is most clearly apparent if the agent is forced to trade all units at the same bid or ask price. Then in order to trade an extra unit a seller will have to lower the price at which he offers all units and conversely a buyer will have to raise the price he bids for all units. But this is not the only way of extending the model and, in fact, there are many ways that one could go about it, with different results.

Perhaps the simplest way of extending the model is to adapt the behavioral rules used in this chapter by representing each agent as a collection of behavioral rules, one for each unit of the good traded. Suppose, for simplicity, that there is a single seller and a single buyer. The seller is assumed to have

N different units of the good and the buyer is assumed to have N different uses for the good. Then the seller's preferences are represented by a sequence of valuations $u = (u_1, ..., u_N)$, where $0 < u_1 < u_2 < ... < u_N$ and the buyer's preferences are represented by a sequence of valuations $v = (v_1, ..., v_N)$, where $v_1 > v_2 > ... > v_N > 0$. The interpretation here is that u_i is the valuation the seller places on unit i of the good and v_j is the valuation the buyer places on a unit of the good applied to use j. Notice that these are not marginal valuations, the value of the n-th unit of the good bought and sold. Rather these are valuations that attach to identifiable units of the good or identifiable uses of the goods. This requires us to assume that the seller regards different units of the good as somehow distinct, independently of the number sold and that the buyer regards different units of the good as identical, but applicable to distinct uses in which they have different valuations, independently of the number of units purchased.

The seller chooses a sequence of ask prices $a = (a_1, ..., a_N)$ where a_i denotes the price asked for unit i and the buyer chooses a sequence of bid prices $b = (b_1, ..., b_N)$, where b_j denotes the price bid for a unit that will be applied to use j. The bids and asks are then matched by the profit-maximizing market maker.

The crucial question is how the bids and asks are chosen. Because of the special assumptions we have placed on preferences, it is possible to assign bid and ask prices to goods independently of the number of goods bought and sold. This means that the prices for different units/uses can be varied independently. In fact, it is possible to regard the seller as a collection of N sellers, each with a different valuation, and the buyer as a collection of N buyers, each with a different valuation. So one way to model the behavior of a buyer or seller is to adopt the behavioral rule used previously by choosing each price to maximize the surplus realized on that unit of the good. Then the previous analysis applies and we get a convergence theorem even with two agents in the market.

As a thought experiment, to see how one might use the results developed in this chapter to understand the behavior of a market in which individuals trade multiple units, this exercise has some value. However, it is not entirely convincing for a number of reasons.

First, an agent is not necessarily choosing a better strategy when he adjusts the ask price for unit i or the bid price for unit j. The reason is that he is competing with himself and increasing the surplus earned on one unit of the good may decrease the surplus earned on another. In fact, he may end up worse off as a result of changing his strategy in this way.

There are two possible responses to this criticism. One response is that, if the object is to model bounded rationality, it may not be such a bad thing that the behavioral rules are somewhat incoherent. Building in more rationality does not necessarily make the model more realistic. Another response is that, if there were many buyers and sellers in the market, the probability that a change in an agent's price for one unit would affect his surplus on another unit would be small and hence might pass unnoticed.

Another, possibly more serious, criticism is that by assuming agents choose the prices for different units independently, we are building into the model an unrealistic amount of competition. There is no possibility for the agent to discover that by varying the prices of different units simultaneously he can exert market power and improve his welfare as a result. In fact, by using the surplus earned on a single unit as the criterion for changing the strategy, we are not even allowing the agent to take into account how his objective function changes as a result of a change in strategy.

Finally, the specification of the preferences is artificial. It does not allow for the usual, natural interpretation that u_i is the seller's valuation of the i-th unit sold and that v_j is the buyer's valuation of the j-th unit bought. If utility is a function of the number of units bought or sold, we cannot specify a strategy for unit i or unit j independently of the other units.

Suppose then that we take the more normal interpretation where each additional unit sold has a higher marginal disutility (cost) and each additional unit bought has a smaller marginal utility. As before, we can illustrate the model by assuming a single buyer and a single seller who choose ask prices $a = (a_1, ..., a_N)$ and bid prices $b = (b_1, ..., b_N)$, respectively. We may want to restrict strategies so that obviously dominated strategies are ruled out, but that is not important as long as the set of admissible strategies is bounded and contains the usual strategies.

The market-maker is, as usual, assumed to choose trades to maximize profits, but here his allocation will have to respect the buyer's and seller's ordering of the bids. If he chooses to trade n units, the profit will be

$$\sum_{j=1}^{n} b_j - \sum_{i=1}^{n} a_i.$$

This is because a seller cannot trade the n-th unit unless he is also trading the first, second, ... and $(n-1)$-th units as well.

At each date, one of the agents is chosen at random to change his strategy. He randomly selects one from the admissible set. If it is better than the current strategy, he adopts it; otherwise he stays with the current strategy.

There are various ways in which search for a better strategy could be structured. An agent could vary the price of one unit at a time or simultaneously change the prices of all units. These modeling choices have important effects on the behavior of the system. Consider first the case where agents vary one price at a time. Suppose that strategies are limited so that $u_i \leq a_i \leq M$ and $0 \leq b_j \leq v_j$. Then the surplus from trade is always non-negative and the agent will never want to restrict his trade, as opposed to carrying on the same trade at more favorable prices. If agents adjust the price for one unit at a time, the model will have the property that once the maximum volume of trade is achieved, it stays the same forever. The reason is that, on the one hand, an agent will never choose to reduce his trade and, on the other, he can only increase his trade by stealing a unit from an agent on the same side of the market. This suggests that the analysis of the model in this chapter could be extended to show convergence if prices are changed one unit at a time.

Would the equilibrium prices converge to the competitive interval? That is not so clear, but the similarity to Bertrand competition suggests that it might be possible to prove a competitive limit theorem, even with a finite number of agents.

The property that the maximum volume of trade, once achieved, is maintained forever depends on the assumption that agents change one price at a time. If agents can change several prices simultaneously the model does not have this property. For example, it may well be advantageous for an agent to reduce his trade if at the same time he gains enough by changing the prices of infra-marginal units. The maximum volume property played an important role in the analysis of convergence in this chapter. If it no longer holds when multiple units are traded and prices on different units are changed simultaneously, this suggests that the analysis of this case will be much more difficult.

If convergence to something holds, will that "something" be perfectly competitive? The fact that agents recognize the price-quantity tradeoff suggests that simultaneous price changing will lead to imperfect competition. In that case, large numbers of buyers and sellers will be needed to achieve a perfectly competitive outcome.

An important special case of simultaneous price changing is one in which the agents must quote a single price for all units traded, possibly accompanied by a limit on the number of units to be traded. A limit on trades is needed because otherwise the agent runs the risk of cornering the market at a disadvantageous price. The interest of this case is that it makes the link between the price quoted and the quantity traded quite tight and so emphasizes the role of imperfect competition. Again, large numbers may be needed to achieve the perfectly competitive outcome.

3.4.3 Multiple Goods

The restriction of the analysis to a single market (i.e., the market for a single good) is just as limiting as the restriction to trading a single unit of the good. If each buyer and seller operates in only one market, then the extension to multiple goods adds nothing to the story developed so far: an economy with many markets is just a collection of single-market economies that do not interact in any meaningful way. So to make the extension interesting, one want agents to buy and sell in more than one market (buy and sell more than one good).

It is not difficult to extend the definition of the model to allow for more than one good, especially if we retain the assumption of quasi-linear utility. For simplicity, suppose that agents can trade at most one unit of each good, but there is a finite number of goods indexed by $h = 1, ..., \ell$ in addition to the numeraire (money). There is a finite number of sellers $i = 1, ..., N$ each of whom sells a vector $x_i \in \{0, 1\}^N$ of goods at the ask prices $a_i \in \mathbf{R}^N_+$ and receives a payoff $a_i \cdot x_i - u_i(x_i)$, where $-u_i(x_i)$ is the disutility of giving up x_i. Similarly, there is a finite number of buyers $j = 1, ..., N$ each of whom purchases a vector of goods $x_j \in \{0, 1\}^N$ at the bid prices $b_j \in \mathbf{R}^N_+$ and receives a payoff of $v_j(x_j) - b_j \cdot x_j$.

There is no difficulty in adapting the behavioral rules to the multi-good case, because there is a price for each different good and the matching of demands and supplies can be carried out for each market independently of the others. Thus, one could choose an agent at random to change his strategy and assume that he randomly chooses a good and randomly chooses a new bid/ask price for that good. If the new price increases the agent's payoff it would be adopted; otherwise the current price would be maintained. This specification has the advantage of simplicity and allows us to extend the analysis of the single-good model in a straightforward way. However, it has the unattractive feature that certain improvements that can only be achieved by changing two prices at once are ruled out.

An example will make this clear. Suppose there are two goods, A and B, and three buyers A, B, and C. Buyer A likes good A, buyer B likes good B, and buyer C likes goods A and B. More precisely, assume that $v_j(0) = 0$ for all j and that

$$\begin{aligned} v_A(0,1) &= 0; v_A(1,0) = v_A(1,1) = 2; \\ v_B(1,0) &= 0; v_B(0,1) = v_B(1,1) = 2; \\ v_C(1,0) &= v_C(0,1) = 1; v_C(1,1) = 6. \end{aligned}$$

So buyer A will pay up to 2 for one unit of A; buyer B will pay up to 2 for one unit of B; and buyer C will pay up to 1 for a unit of either A or B and will pay up to 6 for a bundle consisting of a unit of both A and B. Thus, C has potentially the highest valuation for both goods, but only if he can buy one unit of each. If he can only get one good, buyers A and B each have higher valuations in their respective markets.

For simplicity, assume that there is a single seller with one unit of each good that he values at 0.

Now consider a situation in which A and B are both buying one unit of the good at a price of $1 < p < 2$ and C is bidding less than p for both goods. If C has the chance to change his bid for one of the goods, he will not find it profitable to bid more than p so there is no way that he can become a purchaser. If the seller is asking p then we have an equilibrium. On the other hand, if C could alter both bid prices simultaneously, he could capture the marginal bid on both goods from buyers A and B.

This suggests that agents should be allowed to search for better prices of all goods simultaneously. The analysis of this model will not be so easy, however. Because the marginal valuation of a good depends on the other goods being traded, an agent may well find himself engaged in unrewarding trade. As a result, the process defined by this rule does not have the property of maximizing the volume of trade. An example will make this clear. Suppose that C has the marginal bid for both goods, say, $b_C = (1, 1.5)$. Then buyer A gets a chance to change his strategy and by bidding 1.5 for good A he takes the marginal bid away from C. Then C finds himself paying 1.5 for a unit of good B which is not individually rational. So if C next gets a chance to change his strategy, he may end up choosing not to trade, that is, offering a price below the marginal bid for good B. If the other buyers are also offering low prices for good B there may be no trade in good B. Thus, with multiple goods the volume of trade can fall, unlike the single-good case.

The multi-good model has a tradeoff between price and quantity, in the sense that a trade in one market may influence the price that has to be paid in other. In the example above, buyer C's marginal willingness to pay increases when he consumes a unit of the other good. The opposite could be the case. Suppose that C's preferences satisfy:

$$v_C(1,0) = v_C(0,1) = 1; v_C(1,1) = 1.5.$$

Then C is willing to pay up to 1 for a unit of good A if he consumes no B but would only pay $1.5 - b$ if he were already purchasing one unit of B at a price of b. If buyers A and B trade both goods, they might discover by trial and error that by letting C have a unit of good B they reduce his competition for good A, thus reducing the price they have to pay for it.

An interesting question is whether the interaction of the two markets leads to a kind of imperfect competition, as did the price-quantity tradeoff in the multi-unit, single-good case.

3.4.4 From Partial Equilibrium to General Equilibrium

I have been writing about models with multiple goods as if they were models with multiple markets; but the 'markets' in these models are not really separate and distinct. We can say that there is a market for each good, but this is just a figure or speech. Agents simultaneously trade and quote prices for all of the goods. Since every agent is simultaneously participating in every one of these 'markets', we could equally well say that there is a single market in which all the goods are traded.

Likewise, the models discussed in the preceding chapters sometimes assume a single good (plus money) and sometimes a finite but arbitrary number of goods. We could interpret the former as models of a single market (partial equilibrium models) and the latter as models of an economy with many markets (general equilibrium models), but there is no fundamental difference between them. The specification of behavior and equilibrium are essentially the same. Thus, in the model studied in Chapter 1, all the agents are being matched with each other, presumably in a single location, and when they are matched they offer to trade bundles of commodities rather than individual commodities. There is no sense in which different commodities are being traded on different markets.

If the only difference between a model of a single market and a model of an economy with many markets is the dimension of the commodity space, then it follows that equilibrium in an economy is achieved in the same way as in a single market. This contradicts a long-standing tradition in economic thought, beginning with Keynes if not earlier, that holds the determination of equilibrium in an economy is fundamentally different from the determination of equilibrum in a Marshallian market.

I would argue that although the theory developed here pretends to be a theory of general equilibrium, it does not take seriously the distinction between partial and general equilibrium. A more convincing model of general equilibrium might differ from this one in many ways. One important way in which it ought to differ is in recognizing that markets are distinct and that not all individuals participate in all markets at a given time. This in turn implies that the problem of achieving an efficient allocation of resources is more difficult than the models developed here suggest.