

# Chapter 2

## Continuity and Anonymity

Chapter 1 presents a theory of dynamic matching and bargaining games (DMBG) with a finite number of agents. This theory is more difficult than the corresponding theory for a continuum economy, chiefly because of the strategic problems that arise in finite games. The theory of bargaining deals with small numbers of players, usually only two, and strategic issues arising in these sorts of bargaining problems have been exhaustively studied. Bargaining in markets with a large but finite number of agents is still relatively unexplored territory, however. One of the few papers dealing with this topic is the innovative and imaginative paper of Rubinstein and Wolinsky (1990). It is well worth reviewing the results in this paper because they reinforce the lessons drawn from Chapter 1 about the importance of Markov strategies, anonymity, and continuity.

### 2.1 Rubinstein and Wolinsky (1990)

Rubinstein and Wolinsky (1990), henceforth RW, contains a rich array of models, illustrating the importance of different informational and institutional assumptions in the analysis of DMBG. As a benchmark, RW begins with a model in which the matching process is exogenous. There are  $S$  sellers and  $B > S$  buyers in the market. Each seller has one unit of an indivisible good which is worth nothing to him. Each buyer wants to buy at most one unit of the indivisible good, which is worth one dollar to him. Time is divided into discrete dates  $t = 1, 2, \dots$ . At each date, the agents are randomly matched in pairs consisting of one seller and one buyer (each feasible con-

figuration of pairwise matches has equal probability). One member of the pair is randomly chosen to be the proposer and the other is the responder. Each member has probability  $1/2$  of being chosen as proposer. The proposer chooses a price  $p \in [0, 1]$ . The responder accepts or rejects the offer.

At each date, the agents are randomly rematched. Unmatched buyers are forced to remain inactive for the period.

When agreement is reached on a price  $p$ , the buyer receives a payoff  $1 - p$  and the seller receives a payoff  $p$ . There is no discounting. Agents exit the market as soon as an agreement is reached.

Agents have complete information about the past play of the game, but at the moment when they choose their actions they do not know what actions are being simultaneously chosen by other agents.

The central result that RW obtains for this model is the following theorem.

**Theorem 1** *For every price  $p^*$  between 0 and 1 and for every one to one function  $\beta$  from the set of sellers to the set of buyers there exists a sequential equilibrium in which seller  $s$  sells his unit to buyer  $\beta(s)$  for a price of  $p^*$ .*

In other words, a continuum of outcomes can be supported by sequential equilibria of the game (sequential equilibria are needed here because it is a game of imperfect information). The intuition of the proof was described in Chapter 0 for the case where there is one seller  $S = 1$  and many buyers  $b = 1, \dots, B$ . One buyer  $b^*$  is identified as the intended recipient of the good at the price  $p^*$ . The equilibrium strategies require the seller to offer the good at a price of  $p = 1$  whenever he is the proposer and is matched with a buyer  $b \neq b^*$ . Every buyer  $b \neq b^*$  rejects the offer. Whenever buyer  $b \neq b^*$  is the proposer he offers to buy the good at a price of  $p = 0$  and the seller rejects. When the seller meets buyer  $b^*$ , whichever is chosen as the proposer offers a price  $p = p^*$  and the responder accepts. These strategies clearly produce the required outcome and the payoffs of the seller and buyer  $b^*$  are  $p^*$  and  $1 - p^*$ , respectively, because there is no discounting.

To prevent a deviation, we make use of the following subgame equilibrium strategies. Suppose that the seller has deviated by proposing a price  $p \neq p^*$ . The responder rejects this offer and the game then moves into a subgame in which the formerly rejecting buyer  $b^{**}$  becomes the intended recipient of the good and the selling price becomes  $p^{**} = 0$ . The strategies are the same as those given earlier with the price  $p^{**}$  in place of  $p^*$  and the buyer  $b^{**}$  in place of buyer  $b^*$ .

If one of the buyers deviates by offering a price  $p \neq p^*$  then the seller rejects, another buyer is chosen to be the intended recipient  $b^{**}$  and the price at which the unit is traded changes to  $p^{**} = 1$ . The strategies are defined in the usual way.

Deviations from these punishment strategies can be treated in an exactly similar way.

It is easy to see that the initial deviator is no better off and may be worse off in the punishment phase that is triggered by his deviation. Furthermore, the responder is always at least as well off following the punishment as he would be if he accepted the proposed deviation. It can also be seen, by inspection of the strategies, that a responder who is called upon to reject an offer can never be better off accepting it. Also, a responder who is supposed to accept and rejects instead does not make himself better off if we assume that the game continues as before after a deviating rejection.

These strategies are ultimately quite complicated, in the sense that there is no limit to the number of potential deviations and each additional deviation requires a tailor-made response that makes the play of the game more complicated. As RW point out, one can think of this construction as requiring a large amount of information for the players to execute the equilibrium strategies. For the purposes of comparison, they consider an alternative model in which the amount of information available to the agents is limited. Specifically, the game satisfies the following **anonymity** assumption:

*At the beginning of each date  $t$ , all that the buyers and sellers know about the previous play of the game is the number of buyers  $B_t$  and the number of sellers  $S_t$  remaining in the game.*

Under this assumption, the proposer's strategy is a function of the number of buyers and sellers  $(B_t, S_t)$  and the date  $t$ . The responder's strategy is a function of the numbers of agents  $(B_t, S_t)$ , the date  $t$  and the proposal  $p$ . RW show that, under the anonymity assumption, the only equilibrium outcome is the competitive one.

**Theorem 2** *If each player's information consists only of  $B_t, S_t$ , and  $t$ , the unique sequential equilibrium outcome is such that the good is sold for a price of 1.*

Two points are worth noting. First, the anonymity assumption is identical to the strong Markov assumption, which requires strategies to be functions

of a minimal set of payoff-relevant variables at each date: strategies are not allowed to be conditioned on variables that do not directly affect the future payoffs of the game.<sup>1</sup> Secondly, the assumption of anonymity has the immediate effect of preventing the agents from punishing a deviator, because deviations are not remembered after they occur.

To see how anonymity works, suppose there is one seller and two buyers. The game stops as soon as trade occurs, so  $(B_t, S_t) \equiv (2, 1)$  as long as the game continues. The anonymity assumption implies that the proposals by each agent depend at most on the date  $t$  and the responses depend at most on the date  $t$  and the proposal  $p$ . Of course, different agents can have different strategies.

Let  $1 - \alpha$  be the infimum of the seller's payoff over all subgame perfect equilibria of the game (subgame perfect equilibrium suffices here because there is only one match per period). Because there are two buyers, they cannot both get the good and the sum of their payoffs at any point in the game must be less than or equal to  $\alpha \leq 1$ . Then one of the buyers must be expecting to get no more than  $\alpha/2$  and if the seller gets a chance to make a proposal to this buyer he can ask for anything up to  $1 - \alpha/2$  and be sure of getting it. Call this buyer the "Loser". The existence of a Loser puts a lower bound on the seller's payoff. With probability  $1/4$  the seller is matched with the Loser and the seller is chosen to make the proposal. In that case the seller gets  $1 - \alpha/2$ . With probability  $3/4$  he cannot get less than  $1 - \alpha$ . So his payoff must be at least

$$\frac{1}{4}(1 - \alpha/2) + \frac{3}{4}((1 - \alpha)) = 1 - \frac{7}{8}\alpha \geq 1 - \alpha.$$

The last inequality is strict, which contradicts the definition of  $1 - \alpha$  as the infimum of the seller's payoffs, unless  $\alpha = 0$ . But  $\alpha = 0$  means that the good is always traded at a price  $p = 1$ .

The crucial point in this argument is that each of the buyers has a future payoff that is independent of the offers made by the seller in the current period. In the heuristic proof of Theorem 1, any deviation by the seller was "remembered" by other players and resulted in a reward for the buyer who rejected the deviating proposal. In the anonymous game, by contrast, the buyer's reservation utility is independent of the proposal and this fact allows

---

<sup>1</sup>In Chapter 1 the Markov property was used in a weaker sense: equilibrium strategies were allowed to depend on the date as well as a minimal set of payoff-relevant variables.

the seller to exploit the buyer's weakness, demanding the maximum price that is consistent with the buyer's reservation level.

The absence of discounting is also important. What makes Theorem 1 "work" is the creation of *relationships* among buyers and sellers. At any moment, each unit of the good is dedicated to a particular buyer at a particular price. Each deviation from the equilibrium strategies leads to the creation of a new relationship to punish the deviator. Because the matching process is random and there may be a large number of buyers and sellers, it can take a long time for the designated buyers and sellers to meet one another and consummate the trade. The agents do not care how long it takes for the designated buyer and seller to find each other because there is no discounting. However, if agents discount the future, these relationships can be quite costly to maintain. By the time trade takes place the surplus generated may be heavily discounted. If a seller (buyer) could trade with the first buyer (seller) he meets, he would be much better off, other things being equal. For this reason discounting makes it suboptimal to carry out the punishment strategies that are needed to support the equilibria in Theorem 1. In fact, discounting eliminates large numbers of sequential equilibria altogether. To see this, suppose the model is amended to allow for a common discount factor  $0 < \delta < 1$ . Then it can be shown that there is a unique equilibrium in the case of a single seller. Let  $x(B)$  and  $y(B)$  be defined as the  $(x, y)$  solution of the system of equations

$$\begin{aligned} y &= \frac{\delta(x + y)}{2} \\ 1 - x &= \frac{\delta(1 - x + 1 - y)}{2B}. \end{aligned}$$

**Theorem 3** *Suppose that  $S = 1$  and that all agents discount the future using the common factor  $0 < \delta < 1$ . Then there is a unique subgame perfect equilibrium in which trade takes place immediately and the price is  $x(B)$  or  $y(B)$  depending on whether the buyer or seller was chosen to propose.*

Interestingly, this result has not been established for  $S > 1$ .<sup>2</sup>

---

<sup>2</sup>Suppose that there were a unique equilibrium for  $S > 1$ . What would it look like? If all trade occurs at the first date, then we can use Theorem 3 to calculate the equilibrium payoffs. If a buyer and seller fail to reach agreement at date  $t$  then they anticipate that there will be exactly one seller still active at date  $t + 1$  and the payoffs in that continuation game

Note that for fixed values of  $B$  and  $0 < \delta < 1$ , the surplus is shared between the two sides of the market; but as  $B \rightarrow \infty$  the equilibrium price converges to 1. A large number of buyers makes the equilibrium of Theorem 1, in which one buyer is designated as the good's recipient, extremely inefficient. The exorbitant cost of maintaining relationships in the limit leads to the same outcome as the assumption of anonymity in Theorem 2. An important point to note, however, is that letting  $B \rightarrow \infty$  is increasing the number of buyers relative to the number of sellers, since we are holding  $S = 1$ . This is not the same thing as allowing the size of the market to increase by replication, for example.

Less obvious, but also true, is the fact that as  $\delta \rightarrow 1$ ,  $x(\delta, B) \rightarrow 1$  and  $y(\delta, B) \rightarrow 1$  so that again the seller obtains the maximal price of 1. There is a simple intuition for this result. When  $\delta < 1$  the buyers have some bargaining power because they can threaten to delay agreement. This allows them to "hold up" the seller and get some of the surplus. When  $\delta \rightarrow 1$  this power disappears and the seller is able to exploit fully his monopoly of the good, specifically the power to determine which buyer gets the good.

We can interpret Theorem 3 as showing that the multiple equilibria of Theorem 1 are not robust. There appears to be a close analogy with the role of discounting in the Stahl-Rubinstein alternating offers model. In the alternating-offers game with no discounting, any Pareto-efficient outcome (division of the cake) can be supported by a subgame perfect equilibrium; with a positive amount of discounting, however small, there is a unique subgame perfect equilibrium outcome. So we treat the game with no discounting as a limiting case of the game with discounting and argue that all but one of the equilibrium outcomes in the limit are pathological. In the same way, we can argue that the indeterminacy of equilibrium shown in Theorem 1 is pathological, since it disappears once an arbitrarily small amount of discounting is introduced.

---

will be uniquely determined by Theorem 3. Then in the bargaining at date  $t$  we know that the subgame perfect equilibrium payoffs will be determined in the usual way: the proposal will offer a price that makes the responder indifferent between accepting that price and rejecting it in favor of the payoff from the continuation game. In this equilibrium, the payoffs are determined not by the number of buyers and sellers but rather by the payoffs in the continuation game with  $B - S$  buyers and one seller. Whether this is the only equilibrium is not known. If this is the unique equilibrium, it suggests once again that the effect of large numbers is strongly dependent on the details of the extensive-form game. No matter how large  $B$  and  $S$  are, it is the difference  $B - S$  that appears to be critical in determining the equilibrium payoffs.

RW have a clever comeback to that charge. They point out that the power of discounting to increase the cost of maintaining relationships depends on the assumption that both the formation and the termination of matches are random and exogenous. If the choice of partner were *endogenous*, it ought to be possible to show the existence of multiple equilibria in the presence of discounting. To support this claim, they consider a model in which there is a single seller and the seller can choose, at the beginning of each period, which buyer he wants to bargain with in that period.

**Theorem 4** *If  $S = 1$  and the seller can choose in each period the buyer with whom he wishes to bargain, there is a continuum of subgame perfect equilibrium outcomes: for each buyer  $b$  and each price  $x(1) \leq p \leq 1$ , where  $x(B)$  is defined in Theorem 3, there is a subgame perfect equilibrium in which buyer  $b$  receives the good and the price is either  $p$  or  $\delta p / (2 - \delta)$ , according to whether the seller or buyer  $b$  is the proposer in the first meeting between them.*

Thus, indeterminacy survives discounting.

The indeterminacy found in this model has a family resemblance to an indeterminacy that arises in the Rubinstein-Wolinsky (1985) model described in Chapter 0. In the original paper, it is assumed that a pair of agents separates whenever one of them is matched with a new agent. This behavior is optimal, given that each agent expects all the others to follow the same rule. Since all buyers are identical and all sellers are identical, an agent is indifferent between switching to a new partner and remaining with the old, and if there is a positive probability that his current partner will go off with someone new, the agent strictly prefers to play it safe by hooking up with a new partner himself. However, by a similar reasoning, it would also be optimal for an agent to remain with his current partner and ignore the new matches, if that is what he believes all other agents will do. In that case, two agents, once matched, will stay together until an agreement is reached. The existence of a large number of buyers and sellers becomes immaterial and the equilibrium outcome will be the same as in the two-person bargaining game.

The same two kinds of equilibria exist in the voluntary matching game described in Theorem 4. In one, the seller bargains with a single buyer, ignoring the existence of the other buyers. This equilibrium sets a lower bound to the equilibrium price, namely  $x(1)$ , the price determined in the two-person bargaining game. In the other, competition among the buyers

is allowed full force and the equilibrium price is  $p = 1$ . Once we have two equilibria in the model, it is easy to support more. Other equilibria can be supported by making the decision to separate depend on the agents' behavior.

Although there exists a continuum of equilibrium outcomes in the voluntary-matching model (Theorem 4), the degree of indeterminacy depends on the number of buyers  $B$  and the discount factor  $\delta$ . Just as in the case of exogenous matching with discounting (Theorem 3), the interval of equilibrium prices collapses to  $p = 1$  if either (i) the number of buyers grows unboundedly large, or (ii) the discount factor approaches 1.

In the end, one could use these results to argue either against or in favor of the competitive outcome, depending on which set of assumptions one finds most attractive. The value of the results is really that it shows the importance of the fine details of the model in determining the outcome. Further it clarifies the role of anonymity, continuity, and the Markov property in supporting competitive outcomes.

- Anonymity and other informational restrictions, which RW treat in a way that is equivalent to the Markov property, are used to rule out complex strategies that are required to support indeterminacy in models with large but finite numbers of agents.
- The particular strategies that are used in the proof of Theorem 1 ensure that a single agent may have a large impact on the subsequent play of the game, even when the number of players is large.
- Discounting restores “uniqueness” when matching is exogenous (Theorem 3), but not when it is endogenous (Theorem 4). However, in the voluntary-matching game with discounting, non-Markov strategies are required to support indeterminacy.

It is perhaps also worth noting that another factor that makes it difficult to interpret these results as supporting or contradicting the competitive theory is the fact that the competitive outcome is the same as the monopolistic outcome.

## 2.2 Bounded Rationality and Uniqueness

In the RW game, strategies are sequences of functions mapping histories into actions. If we think of the complexity of a strategy as being measured by

the number of variables on which an agent's action can be conditioned, then the strategies available in some versions of the RW game are unboundedly complex. More precisely, since there is no bound on the length of a history, there is no bound on the complexity of the strategy that may be adopted. In this section we explore a restriction on agents' strategies that is less restrictive than the Markov property, but nonetheless allows for some bound on the ability of the agents to condition their behavior on the history of the game. Whereas RW interpret their restrictions on strategies in terms of the information available to agents, the restrictions considered here are interpreted as a form of **bounded rationality**: there is a limit on the amount of information about past play of the game that an agent can retain and process in determining the future play of the game. For the purpose of this discussion, I retain the assumptions that appear to be most amenable to indeterminacy, namely, voluntary matching and no discounting. Also, only the case with a single seller and finite number of buyers is considered.

### 2.2.1 Bounded memory states

The Markov assumption eliminates the use of any information from the past to condition the future play of the game. In fact, the assumption that strategies are memoryless is the most restrictive assumption that is consistent with existence of an equilibrium. At the other extreme, the entire history of the game can be remembered and used to condition the future play of the game. In this section, we take an intermediate course, by allowing agents to remember a limited amount of what happened in the past. By placing a bound on the agents' memory we put a bound on the complexity of their strategies.

An agent's memory is represented by a set of **memory states**, each of which represents an admissible amount of information that can be retained from the past. Let  $S$  denote the set of memory states and note that  $S$  is common to all agents. The seller makes offers and chooses the buyer to whom the offer is made at each date. A strategy for the seller is a function  $f : S \rightarrow I \times [0, 1]$  that maps a memory state  $s$  into an ordered pair  $(i, p)$ , with the interpretation that the seller offers buyer  $i$  the object for a price of  $p$ . A strategy for a buyer  $i$  is a function  $g_i : S \times [0, 1] \rightarrow \{Y, N\}$  that maps an ordered pair  $(s, p)$ , representing an offer of the object to agent  $i$  at price  $p$ , into a response  $Y$  or  $N$ . These strategies apply only if the object has not already been sold, of course. Given a memory state  $s$ , the strategy profile  $(f, g) = (f, g_1, \dots, g_n)$  determines a unique action  $a = (i, p, r)$  according to

the equations

$$(i, p) = f(s) \tag{2.1}$$

and

$$r = g_i(s, p) \tag{2.2}$$

The memory state evolves according to an exogenously given rule  $\phi$ . This rule represents the agent's ability to encode new information that he receives in the language of the memory states. Formally,  $\phi$  maps the current memory state and the actions chosen in the current period into the future memory state. If  $s_t$  is the current memory state,  $a_t$  the joint action at date  $t$ , and  $s_{t+1}$  the memory state at date  $t + 1$ , then the law of motion is written

$$s_{t+1} = \phi(s_t, a_t) \tag{2.3}$$

where  $a_t = (i_t, p_t, r_t)$ ,  $(i_t, p_t) = f(s_t)$  and  $r_t = g_{i_t}(s_t, p_t)$ . For any strategy profile  $(f, g)$  and initial state  $s_0$ , there is a unique path determined by the equations (2.1), (2.2), and (2.3). With this set of strategies and this path we can associate a unique payoff function for the players. Let  $U(f, g, s_0)$  (resp.  $V_i(f, g, s_0)$ ) denote the seller's (resp. buyer  $i$ 's) payoff from the strategy profile  $(f, g)$  and initial state  $s_0$ .

If the set of memory states is rich enough, this way of describing strategies does not restrict the game at all. For example, if  $S$  is the set of all possible histories  $H$  and  $\phi : H \times A \rightarrow H$  is defined by

$$\phi(h, a) \equiv (h, a),$$

then the game in which strategies are defined on memory states is strategically equivalent to the RW game in which strategies are defined on complete histories. The point of this formalism is to allow for flexibility in choosing restrictions on the set of strategies. For example, if we take the set  $S$  to be a singleton, then we are restricted to the Markov strategies.<sup>3</sup> Assuming  $S$  to be a finite set allows for a richer set of strategies, but one that limits complexity when  $n$  is large. Restricting the set  $S$  to be finite, we can show that there is an essentially unique outcome, in which the seller gets all the surplus, when the number of players is sufficiently large.

---

<sup>3</sup>The game continues only if there has been no trade, so there is only one payoff-relevant state in which the game continues.

**Theorem 5** *If  $(f, g)$  is a subgame perfect equilibrium of the game with a fixed set of memory states  $S$  and  $|S| < n$ , then the seller offers the object to one of the buyers at a price  $p = 1$  and the buyer accepts.*

**Proof.** The equilibrium payoff for the seller (resp. buyer  $i$ ) in state  $s$  is denoted by  $u(s) = U(f, g, s)$  (resp.  $v_i(s) = V_i(f, g, s)$ ). It depends only on the state  $s$  because, for a given strategy profile  $(f, g)$ , the future play of the game is a function of the current state. Feasibility requires that

$$0 \leq \sum_{i=1}^n v_i(s) \leq 1 \quad (2.4)$$

for all states  $s$ . This implies that there exists a buyer  $i$  such that

$$v_i(s) \leq \frac{1}{2}, \forall s \quad (2.5)$$

Otherwise, for each buyer  $i$  there exists a state  $s_i$  such that  $v_i(s_i) > 1/2$ . Since  $|S| < n$ ,  $s_i = s_j$  for some pair of agents  $i \neq j$ . In other words, there exists a state  $s$  such that  $v_i(s) > 1/2$  for at least two buyers, contradicting the feasibility condition (2.4). This contradiction proves that (2.5) must hold for some  $i = k$ , say.

The condition (2.5) allows us to place a lower bound on the equilibrium payoff of the seller. Given any state  $s$ , the seller knows that agent  $k$  cannot expect a payoff higher than  $1/2$  from the continuation of the game. Therefore agent  $k$  should accept any price offer  $p < 1/2$ . But this means that the seller can get a payoff of  $1/2 - \varepsilon$  by offering a price  $p = 1/2 - \varepsilon$  to buyer  $k$ . Since  $\varepsilon$  is arbitrary, the seller must get  $1/2$  in equilibrium:

$$u(s) \geq \frac{1}{2}, \forall s. \quad (2.6)$$

Condition (2.6) and feasibility imply a new bound on the sum of buyers' payoffs

$$0 \leq \sum_{i=1}^n v_i(s) \leq \frac{1}{2}, \forall s$$

and this implies by the previous argument that for some agent  $i$

$$v_i(s) \leq \frac{1}{4}, \forall s.$$

Continuing in this way, we can see that, in the limit, we must have  $v_i(s) = 0$  for all  $s$  and  $i$  and  $u(s) = 1$  for all  $i$ . ■

The proof of Theorem 5 uses induction. An alternative proof, which is less heuristic but more efficient, does not require induction: if we let

$$\alpha = \min_i \max_s \{v_i(s)\}$$

then we can see that  $u(s) \geq 1 - \alpha$  and, by the usual argument, this implies that  $\alpha \leq \alpha/2$  and hence  $\alpha = 0$ . By choosing to deal with the weakest buyer  $i$ , for whom  $\max_s \{v_i(s)\} = 0$ , the seller can guarantee himself the maximum payoff.

It is obvious that trade must take place in equilibrium or the payoffs in Theorem 5 cannot be achieved. However, trade does not have to take place at the first date. Although the agents' strategies do not change over time the state  $s$  can change over time and this allows us to introduce delay in the equilibrium strategies. For example, let  $S = \{1, 2, \dots, T\}$  and  $s_0 = 1$  and suppose that for any action profile  $a$  the evolution of the memory state is

$$\phi(s, a) = \begin{cases} s + 1 & \text{if } s < T \\ s & \text{if } s = T \end{cases} .$$

Then the memory state is essentially counting time and we can use it to code the date at which trade takes place. For example, at any  $s < T$  the seller offers the object at a price  $p = 1$  to buyer  $i$  and buyer  $i$  rejects the offer. At  $s = T$  the seller offers the object at a price  $p = 1$  to buyer  $i$  and  $i$  accepts. Since there is no discounting, both agents are indifferent among the  $T$  dates at which trade could take place.

The argument in the proof of Theorem 5 is similar to the argument in the heuristic proof of Theorem 2. In the anonymous equilibrium, there is a single memory state (the Markov case) and there are two buyers, so the number of buyers is greater than the number of memory states. For the proof of Theorem 5, it is only necessary to extend the argument to an arbitrary finite number of memory states and a (larger) finite number of buyers. What makes both arguments work is the fact that there is only one unit of surplus available to reward all the buyers for punishing the seller. In the proof of Theorem 1, this one unit of surplus is cleverly used to reward the agent who responds to a deviant proposal: by rejecting the deviating offer the responder gets all the surplus, and that is enough to discourage him from accepting any

deviation from the equilibrium path. The assumption that the set of states is less than the number of buyers prevents the set of strategies from being rich enough to implement this scheme of punishment and reward. There will always be one buyer who is a “Loser”, that is, gets less than one unit of surplus no matter what the seller does. The seller can target this buyer and extract the maximal amount of surplus. The Loser is the weak link in the chain. The existence of one Loser ensures that all the buyers are Losers and ultimately they all end up with nothing.

The interpretation of the outcome of this game as “competitive” may not be entirely convincing. It is well known in bargaining games with discounting that allowing one agent to make all the offers typically allows that agent to capture all the surplus. Since the monopolistic outcome is the same as the competitive outcome, the fact that the seller captures all the surplus could be interpreted, quite naturally, as a monopolistic outcome. However, the same outcome can be obtained in a model in which the proposer is chosen randomly. For example, suppose that the seller chooses the buyer  $i$  with whom to bargain and then the proposer is chosen with equal probability to be the seller or buyer  $i$ . Since there is no discounting and the seller will be chosen as proposer infinitely often, the argument above can be used to prove that the seller can force an outcome in which he gets all the surplus.

### 2.2.2 Imperfect memory

The bounded memory assumption clarifies one sense in which bounded rationality, or bounded complexity of strategies, can select the competitive outcome from a large, possibly infinite set of equilibria. At the same time, it is a rather strong assumption. Another approach to this problem does not restrict the number of memory states in relation to the number of players. Instead, it introduces a “small” amount of randomness to mimic the effects of imperfect memory. Specifically, for some fixed value of  $\varepsilon > 0$  we assume that, with probability  $1 - \varepsilon$ , the transition is determined equation (2.3) and with probability  $\varepsilon$  the new state is uniformly distributed on  $S$ :

$$s_{t+1} = \begin{cases} \phi(s_t, a_t) & \text{w. pr. } 1 - \varepsilon + \varepsilon/|S| \\ s \neq \phi(s_t, a_t) & \text{w. pr. } \varepsilon/|S|. \end{cases}$$

The random element in the transition is like a “trembling hand” in the play of the game. It ensures that there is always a small probability that the players

in the continuation game are following the “wrong” strategy, for example, rewarding the wrong player.

This element of bounded rationality in the game leads to the same outcome as before. Intuitively, the small amount of randomness prevents any one buyer from being awarded the entire surplus. This ensures that the seller can guarantee himself a positive amount of the surplus in any state. This in turn reduces the maximum that the buyers as a group can hope to get and the small amount of randomness ensures that any buyer must get less than this new, lower maximum. Then the seller can guarantee himself even more. Pursuing this argument to its logical conclusion shows that, in equilibrium, the seller must get the entire surplus.

**Theorem 6** *If  $(f, g)$  is a subgame perfect equilibrium of the game with imperfect memory and  $n > 2$ , then the seller offers the object to one of the buyers at a price  $p = 1$  at the first date and the buyer accepts.*

**Proof.** As before, let  $(f, g)$  be a fixed subgame perfect equilibrium and put  $u(s) = U(f, g, s)$  and  $v_i(s) = V_i(f, g, s)$  for  $i = 1, \dots, n$ . Let  $1 - \alpha$  be the minimum payoff to the seller, that is,

$$1 - \alpha = \min\{u(s) : s \in S\}.$$

Then feasibility implies that

$$\sum_{i=1}^n v_i(s) \leq \alpha, \forall s. \quad (2.7)$$

From the feasibility constraint (2.7) it follows that, for some buyer  $i$ ,  $v_i(s) \leq \alpha/2$  for at least a fraction  $(n - 1)/n$  of the states. To see this, let

$$S_i = \{s \in S \mid v_i(s) > \alpha/2\}.$$

By the definition of  $S_i$  and the feasibility condition (2.7) it is clear that  $S_i \cap S_j = \emptyset$  for  $i \neq j$  so there must exist an agent  $i$  such that

$$|S_i| \leq \min_{j=1, \dots, n} |S_j| \leq \frac{|S|}{n}.$$

For this agent  $i$ ,  $|S \setminus S_i| \geq \frac{n-1}{n}|S|$  as claimed.

Now consider the expected payoff of agent  $i$  for any memory state  $s_t$  and action  $a_t = (i, p)$ :

$$\begin{aligned} E[v_i(s_{t+1})|(s_t, a_t)] &\leq (1 - \varepsilon)\alpha + \varepsilon \left( \frac{1}{n}\alpha + \frac{n-1}{n} \frac{\alpha}{2} \right) \\ &= \left[ \frac{\varepsilon}{2} \left( \frac{n+1}{n} \right) + (1 - \varepsilon) \right] \alpha \\ &= \alpha - \varepsilon\alpha \left( 1 - \frac{n+1}{2n} \right). \end{aligned}$$

As long as  $n > 2$ , the preceding inequality shows that  $v_i(s) < 3\alpha/4$  for every  $s$  and, in a subgame perfect equilibrium, this means that agent  $i$  must accept any offer  $p < 1 - \alpha(1 - \varepsilon/2)$ . Then the seller can guarantee himself a payoff arbitrarily close to  $1 - \alpha(1 - \varepsilon/2)$ . This contradicts the definition of  $1 - \alpha$  unless  $\alpha = 0$ . ■

The requirements for Theorem 6 are weaker than for Theorem 5: Theorem 6 only requires the existence of what we might call an  $\varepsilon$ -Loser, that is, someone who receives less than  $(1 - \varepsilon)\alpha$  in every state, and  $\varepsilon > 0$  can be arbitrarily small. In the same terminology, Theorem 5 requires the existence of a  $1/2$ -Loser.

The assumptions made in Theorem 6 are in one respect stronger than necessary. The argument never refers to the number of elements in  $S$ . In fact, looking at the details of the proof, it appears that the same argument applies to an arbitrary measurable space  $S$ , for example, suppose that  $S$  is a compact metric space endowed with the Borel sets  $\mathcal{B}(S)$  and a non-atomic probability measure  $P$ .

What both theorems teach us is that it is not necessary to assume completely memoryless strategies in order to rule out all non-competitive behavior. Strategies can be conditioned on a large amount of information about the past, as long as the complexity of the strategies is restricted so as to prevent the kinds of complicated punishments needed to support the continuum of equilibria described in Theorem 1.

Following a different approach, Chatterjee and Sabourian (1998) have analyzed bargaining games between finite automata. The idea of requiring strategies to be played by finite automata (finite state machines) is intended to impose some bound on the complexity of play. The automaton is designed to play a best response, but it has the minimal number of states needed along the equilibrium path. This assumption does not eliminate

non-uniqueness, but it does limit the amount of delay that can occur in equilibrium.

## 2.3 The Limit Principle

Non-cooperative models of oligopoly have long been used to provide a strategic foundation for competitive equilibrium (see discussion in Chapter 0). The earliest and best known approach is based on the well known Cournot model of oligopoly. In this model, there are  $n$  identical firms, each of which simultaneously chooses the quantity of output it wants to produce. An inverse demand function, representing the behavior of the consumers in the market, determines the market-clearing price as a function of the aggregate output. The profit of each firm is determined by the price and the firm's output, so we have a well defined game in which the payoff of each firm is a function of the  $n$ -tuple of quantities chosen. The outcome is described by a Nash equilibrium of this game.

When the number of firms in the market is finite, the static Cournot oligopoly game predicts that competition will be imperfect, and the price will be higher and output lower than in the competitive, price-taking equilibrium. However, as the number of firms becomes unboundedly large, the oligopolistic equilibria converge to the perfectly competitive, price-taking equilibrium. Some strong regularity conditions are required for the proof of this limit theorem and, in an important paper, Roberts (1980) provided a counter-example. Roberts (1980) shows that, under certain conditions, the market demand curve is not invertible. A *selection* from the demand correspondence may be inverted, but the resulting inverse demand curve will not be continuous. Specifically, at some point, an increase in the aggregate quantity supplied may cause a sharp (discontinuous) fall in the market price. The effect of this discontinuity is to prevent the increasing competitive pressure from expanding aggregate output beyond the discontinuity point. When the market is poised at a discontinuity point, producing an extra unit of output causes a loss on inframarginal units that is a higher order of magnitude than the revenue from selling an extra unit. No matter how many firms are in the market, the aggregate output may not increase beyond this discontinuity point.

Another way of looking at the problem is to say that, no matter how many firms and consumers are in the market, each firm continues to have a

non-negligible effect on the market, since the slightest increase in output by one firm pushes the market over the precipice and leads to a sharp fall in the market price.

The Roberts (1980) example illustrates a more general problem. Unless some restrictions are placed on the payoff functions to ensure continuity of the payoffs as functions of the strategy profiles, there is no hope of achieving a limit theorem characterizing the behavior of sequences of Nash equilibria as the number of players increases without bound. This problem was studied in Green (1984)'s work on limits of Nash equilibria in games. The problem posed by Roberts (1980) and studied by Green (1984) may appear to be a purely technical problem, one that can be remedied by assuming that games are sufficiently well behaved. However, a more robust form of discontinuity arises quite naturally in repeated games.

One weakness of the static Cournot model is that it does not capture the possibilities for collusion in a realistic way. In most industries, the competition among firms continues over many periods. This repeated interaction is naturally represented by a repeated game in which the stage game at each period is the static Cournot oligopoly game. Once we allow for repeated interaction, the set of (subgame perfect) equilibria becomes much larger. According to the Folk Theorem for repeated games (see, e.g., Fudenberg and Tirole (1992)), any profile of individually rational payoffs can be supported by a subgame perfect equilibrium, as long as the rate of time preference is sufficiently low. In particular, there will be an equilibrium in which firms collude to share the monopoly profits.

One way of supporting the collusive outcome is by using the "Nash reversion" strategies. Suppose there are  $n$  firms in the market, operating under constant returns to scale. If  $y^m$  is the monopolistic level of output, then each firm is assigned an output of  $y^m/n$  and receives a fraction  $\pi^m/n$  of the monopoly profits  $\pi^m$  each period. However, if any firm deviates from the output level  $y^m/n$ , then every firm reverts in the next period to the Nash-Cournot equilibrium of the one-shot game, that is, to the level  $y^c$ . Playing an equilibrium strategy of the one-shot game in every period is a subgame perfect equilibrium strategy, and since the profits in every period will be lower under this strategy, the punishment on the deviating firm will be large enough to deter the deviation, as long as the discount rate is sufficiently low.

If we replicate this market by increasing the number of firms and consumers, it actually becomes easier to support the monopoly outcome. The monopoly outcome does not change, but the Nash-Cournot equilibrium of

the one-shot game converges to the price-taking equilibrium, so the threat of reverting to the Nash-Cournot equilibrium becomes larger and more effective. Thus, increasing the number of firms does nothing to undermine the possibilities for collusive behavior.

This problem was studied by Edward Green in a prescient paper (Green (1980)). The focus of Green's study was again the Cournot model of oligopoly. Like Roberts, Green was interested in what he called the *Limit Principle*, under which the equilibria of the static Cournot oligopoly game converge to a price-taking competitive equilibrium when the number of firms becomes very large. As we have seen, the Limit Principle fails to hold in the repeated game.

Another way of thinking about the failure of the Limit Principle is as a failure of continuity in the limit: even though the number of firms becomes unboundedly large, so that the single firm becomes negligible relative to the entire market, the actions of a firm are still significant, because the entire market reverts to the Nash strategies when a single firm deviates.

It is also worth noting that the equilibrium strategies do not satisfy the Markov property, because the punishments are based on the history of the game.

This behavior of sequences of replica economies is in marked contrast to the limit economy with a non-atomic continuum of firms. Suppose that instead of replicating the economy we had assumed the existence of a continuum of firms with the same cross-sectional distribution of characteristics as the finite market. (Previously we assumed identical firms with constant returns to scale; now firms are allowed to be heterogeneous). One way of doing this is to replace each firm in the finite market with a non-atomic continuum of firms having identical cost functions. Suppose further that the market game is *anonymous*, in the sense that firms cannot observe the output decisions of other firms; they can only observe the price in the market at each date. Now a single firm in this continuum model has literally no effect on price, because its output has no effect on the total quantity produced. (Each firm has measure zero and so cannot affect aggregate output). As a result, each firm will take the price as given and it is easy to see that the equilibrium at each date will be the standard price-taking or perfectly competitive equilibrium of the static market. The firm does not have to worry about threats of retaliation. It can do whatever it likes without any fear that its deviation will be detected. In fact, this is just the Anti-Folk Theorem for repeated anonymous games (see Rosenthal and Masso (1989)).

This points to the crucial difference between the continuum model and

the large-but-finite model. In the continuum model, the firm has no impact on the equilibrium price. In the large but finite model, no matter how many times the market has been replicated, a single deviation will lead to some change in price and this price change can trigger the punishment phase in which every firm reverts to the Nash-Cournot equilibrium. A single firm has a small direct impact on price, because the firm's output is small relative to the size of the market. However, that small price change has a large strategic impact on the future play of the game and hence on the future market price.

Green's insight is that the collusive equilibrium is unlikely to be robust. In practice, many random and unobservable factors would influence the price in addition to the actions of the firms. It is unlikely that firms in general will have enough information to infer from an observation of the price whether or not some firm has deviated from the collusive equilibrium. If enough noise enters the information channel, firms can cheat on the collusive agreement without fear of detection. In that case, the Limit Principle might be reinstated.

Green (1980) contains a detailed study of the limiting properties of sequences of replicated, anonymous, repeated games in which players can observe noisy public signals that depend on individual actions. Although the market studied is very specific, the framework developed in the paper is very general and applies to other repeated games. Green's results are generalized by Sabourian (1990), which makes use of the same framework and extends the analysis in several directions. Green considers only equilibria supported by Nash reversion strategies and also allows only strategies in which agents' actions depended on public signals. Sabourian's analysis applies to general equilibria, not just equilibria with Nash reversion strategies, and allows strategies to depend on past signals as well as the agent's own past actions.

The next few sections of this chapter are devoted to re-interpreting the analysis of Green and Sabourian in the context of repeated games with limited memory. The objective is to prove a limit theorem analogue of the Anti-Folk Theorem. This theory will apply not just to market games, but to repeated games in general. This treatment is more limited than Green (1980) and Sabourian (1990)—I do not attempt to characterize the limiting continuum game and the arguments are developed slightly differently—and the arguments are more direct and less abstract. The underlying ideas are the same, however. For an elegant treatment, the reader is referred to Sabourian (1990).

The next two sections describe the general framework. In the section after

that I derive a limit theorem and consider the implications for the analysis of dynamic games.

## 2.4 Repeated Games

A finite, normal-form game is denoted by  $\Gamma = (N, X, u)$ , where

- $N = \{1, \dots, n\}$  is the set of players;
- $X_i$  is the (finite) strategy set of player  $i$  and  $X \equiv \times_{i=1}^n X_i$  is the set of strategy profiles;
- $u_i : X \rightarrow \mathbf{R}$  is the payoff function for player  $i$  and  $u \equiv (u_1, \dots, u_n)$  is the payoff function for the game.

The game is extended to allow for mixed strategies in the usual way. Let

$$\Delta(X_i) \equiv \left\{ p_i : X_i \rightarrow \mathbf{R}_+ \mid \sum_{x_i \in X_i} p(x_i) = 1 \right\}$$

denote the set of mixed strategies and  $\Delta(X) \equiv \times_{i=1}^n \Delta(X_i)$ . Then define  $u_i : \Delta(X) \rightarrow \mathbf{R}$  by putting

$$u_i(p) = \sum_{x \in X} P(x) u_i(x)$$

where  $p \equiv (p_1, \dots, p_n)$  and  $P(x) \equiv \times_{i=1}^n p_i(x_i)$ .

A Nash equilibrium of  $\Gamma$  is a strategy profile  $p^*$  such that for any player  $i$

$$u_i(p^*) \geq u_i(p_{-i}^*, p_i), \forall p_i \in \Delta(X_i),$$

where  $p_{-i}^*$  denotes the profile of strategies of players other than  $i$  and  $(p_{-i}^*, p_i) \equiv (p_1^*, \dots, p_{i-1}^*, p_i, p_{i+1}^*, \dots, p_n^*)$ . The set of Nash equilibria is denoted by  $NE(\Gamma)$ .

A repeated game consists of an infinite number of plays of a finite game, with play of the game occurring at a sequence of dates  $t = 1, 2, 3, \dots$ . At each date the players can observe the play of the game in all of the preceding stages and condition their play in the current stage on that information. The payoff in the repeated game is the present value of the payoffs from the sequence of stage games.

Let  $\Gamma^\infty$  denote the countable repetition of the finite stage game  $\Gamma$ . The set of dates at which play occurs is denoted by  $T = \{1, 2, \dots\}$ . The history of the game at date  $t$  is a finite sequence of (pure) strategy profiles,  $\{x_s\}_{s=1}^t$  in  $X$ . Let  $H$  denote the set of histories of the game, including the null history  $\emptyset$ . Then a strategy for player  $i$  is a function  $f_i : H \rightarrow \Delta(X_i)$ , with the interpretation that  $f_i(h) \in \Delta(X_i)$  is player  $i$ 's mixed strategy in the stage game  $\Gamma$  at the information set  $h$ . A profile for the repeated game is a function  $f : H \rightarrow \Delta(X)$ . Let  $F_i$  denote the set of strategies for player  $i$  and let  $F = \times_{i=1}^n F_i$  denote the set of strategy profiles. For each strategy profile  $f$  there is a unique path  $\xi(f) = \{\xi_t(f)\}$ . This path is random if the players choose mixed (behavioral) strategies and has a probability distribution  $\mu_f$ . We can define the payoff function by putting

$$U(f) = E \left[ \sum_{t=1}^{\infty} \delta^{t-1} u(\xi_t(f)) \right] \quad (2.8)$$

for every  $f$ , where  $\delta \in (0, 1)$  is the common discount factor for all players, and the expectations operator  $E[\cdot]$  is with respect to  $\mu_f$ .

A Nash equilibrium of  $\Gamma^\infty$  is a strategy profile  $f^*$  such that for each  $i$  and any strategy  $f_i$

$$U_i(f^*) \geq U(f_{-i}^*, f_i).$$

For a fixed history  $h$  and a strategy  $f$ , we define a new strategy  $f|h$  by putting

$$(f|h)(h') = f(h, h').$$

Then we say that  $f^*$  is a subgame perfect equilibrium (SPE) if  $f^*|h$  is a Nash equilibrium for every  $h$ .

## 2.5 Limited Memory

In this section, I extend the framework of the repeated game to allow for the possibility that players cannot condition their behavior on the precise history of the game. This might be the result of bounded rationality, limited memory, or some other departure from the assumptions of the standard model.

Let  $S$  be a set of states of the game. These states represent what the players can remember about the history of the game. The states do not have any direct impact on the agents strategies or their payoffs. The only function of the states is to allow players to condition their actions on some

limited information about the past play of the game.  $S$  is assumed to be a metric space. The measure space  $(S, \mathcal{B}(S), \lambda)$  is formed by endowing  $S$  with the  $\sigma$ -field  $\mathcal{B}(S)$  consisting of the Borel sets of  $S$  and a  $\sigma$ -additive, positive measure  $\lambda$ . (For definitions and standard results, see any text on probability and measure theory, e.g., Billingsley (1985)).

The evolution of the state is determined by the play of the game through a probability transition function. For each strategy profile  $x$  and state  $s$ , the probability distribution of the state in the next period is denoted by  $\psi(x, s)$ . Then  $\psi$  satisfies the following properties:

- $\psi : X \times S \rightarrow \Delta(S)$ , where  $\Delta(S)$  is the set of probability measures on  $(S, \mathcal{B}(S))$ .
- For any measurable set  $A \in \mathcal{B}(S)$ ,  $\psi(A|x, s)$  is a measurable function of  $(x, s)$ .

The initial state of the game  $s_1$  is assumed to be fixed.

A repeated game with limited memory  $\Gamma_L^\infty(s_1)$  is defined by the repeated game  $\Gamma^\infty$  together with the measure space of states  $(S, \mathcal{B}(S), \lambda)$ , the transition probability function  $\psi$  and the initial state  $s_1$ . Formally, this game is a stochastic game, but it is a degenerate kind of stochastic game because the states do not affect the payoffs or the set of strategies available. They are akin to the “sunspots” in models of correlated equilibria.

By definition, all strategies in a game of limited memory have the strong Markov property: a strategy profile for the stochastic game is a measurable function  $f : S \rightarrow \Delta(X)$ . In other words, the strategy profile chosen in the stage game  $\Gamma$  is a function of the current state of the game only. Together with the transition function  $\psi$  the strategy  $f$  defines a stochastic process  $\xi(f) = \{\xi_t(f)\}_{t=1}^\infty$  and the payoff function  $U$  can be defined in the same way as before by (2.8). A Markov perfect equilibrium (MPE) for the stochastic game  $\Gamma_L^\infty$  is a Markov strategy profile  $f^*$  such that for any date and state  $s$ , the strategy  $f^*$  is a Nash equilibrium of the game  $\Gamma_L^\infty(s)$ .

### 2.5.1 Markov Perfect Equilibria

It may appear that the MPE of the stochastic game are special, but in fact there is no essential loss of generality in focusing on MPE as long as the state space  $S$  and the transition probability  $\psi$  are suitably defined. Recall that

the states do not affect anything of substance in the stage game  $\Gamma$ . They are merely “sunspots”. So we can use them to carry information about the past. If this is done appropriately, then any SPE of the repeated game can be replicated as a MPE of the stochastic game.

**Theorem 7** *Suppose that  $f^*$  is a SPE of the repeated game  $\Gamma^\infty$  and consider the stochastic game  $\Gamma_L^\infty(s_1)$  such that  $S \equiv H$ ,  $s_1 = \emptyset$  and  $F$  is the “deterministic” transition probability function defined by putting*

$$\psi(h'|x, h) = \begin{cases} 1 & \text{if } h' = (h, x) \\ 0 & \text{otherwise,} \end{cases}$$

*for any  $(h', x, h)$ . Then we can define a MPE  $g^*$  of  $\Gamma_L^\infty(s_1)$  by putting  $g^*(h) = f^*(h)$  for every  $h \in H$ .*

**Proof.** The statement of the theorem is almost a proof. It is enough to show that for every state  $h$  the payoff associated with  $g^*$  will be the same as the payoff in the repeated game after the history  $h$ . ■

Note that the particular state space and transition probability used for the theorem are not essential. As long as  $S$  is (at least) countably infinite and the transition probability is an invertible function, Theorem 7 continues to hold. The essential thing is that the structure of the game allows us to encode the history in the set of states.

Theorem 7 shows that the Markov assumption loses its power to reduce the set of equilibria if the state space is appropriately expanded. The point of introducing this result is simply to show that there is no loss of generality in making the Markov assumption in this context.

While Theorem 7 shows that we can replicate the SPE of a repeated game as MPE of a game with limited memory, the replication requires a huge state space. The fact that these strategies are Markov does not mean that they are simple strategies when the state space is so large. In practice, it may seem unlikely that individuals could implement strategies that condition on such a complex memory space. Even if they had the ability to implement such strategies, how would they calculate the equilibrium strategies? Restricting the state space is a natural and attractive way to capture the notion that agents have some bounds on their rationality (ability to perform complex computations).

One way of representing “bounded rationality” is that to assume that  $S$  is finite. This assumption includes as special cases a number of well known

models of bounded rationality. An example introduced by Green (1982) and Rubinstein (1986) is the *finite automaton*, that is, a finite-state machine. The machine representing player  $i$  has a finite number of states, denoted by the set  $S_i$ , and his strategy (output) is a function  $f_i : S_i \rightarrow X_i$ . The player's state evolves according to a rule (transition function)  $g_i : S_i \times X \rightarrow S_i$ . In words, if his state at date  $t$  is  $s_i$  and the action profile is  $x$  then player  $i$ 's state at date  $t + 1$  will be  $s'_i = g_i(s_i, x)$ . Both  $f_i$  and  $g_i$  are considered endogenous, in other words, both are chosen to satisfy the equilibrium conditions. In practice, the set of states is kept as small as possible. The interest of this kind of analysis lies in seeing what kind of equilibria can be supported by very simple strategies.

To see that the finite automaton is a special case of the present framework, simply define the memory states as  $n$ -tuples of individual (machine) states, that is, put  $s = (s_1, \dots, s_n)$  and use the function  $g = (g_1, \dots, g_n)$  to define the (deterministic) transition probability  $\psi$ ,

$$\psi(s'|x, s) = \begin{cases} 1 & \text{if } s' = g(s, x) \\ 0 & \text{otherwise.} \end{cases}$$

The strategy  $f_i$  can be treated as a function of  $s$  as long as we remember that it should depend only on the component  $s_i$ .

Another model of bounded rationality assumes that players have finite memories. The assumption of finite memory is more restrictive than the assumption of a finite memory space  $S$ . An agent with finite memory can only remember what has happened in a finite number of recent periods. This implies that strategies can only depend on the history of a finite number of periods. For example, the strategy of player  $i$  at date  $t$  would be a function  $f_i(x_{t-T}, \dots, x_{t-1})$  of the history of play in the last  $T$  periods. In that case, we put  $S = X \times \dots \times X$  and define the transition probability by putting

$$\psi(s'|x, s) = \begin{cases} 1 & \text{if } s' = (x_{t-T+1}, \dots, x_{t-1}, x) \text{ and } s = (x_{t-T}, \dots, x_{t-1}) \\ 0 & \text{otherwise.} \end{cases}$$

This is clearly just a variant of the framework described earlier, where we put  $S = H$ . An example of finite memory is found in Aumann and Sorin (199?), who study cooperation in two-player coordination games. The players' actions are assumed to depend on a finite history of their opponent's actions. This means that any deviation from equilibrium play is eventually forgotten by the opponent and this allows players to experiment without fear

of permanent punishment. It is important in the Aumann-Sorin setup that players do not condition their actions on the history of their own play. If they did condition on their own past play, they could use the fact that they had punished last period as a signal to punish again this period. In this way, their own past actions could act as a permanent reminder of an opponent's deviation and this would undermine the effect of the finite memory assumption.

These two examples show that the present framework can accommodate familiar models of bounded rationality, but it is more general than these models of individual behavior. A finite memory space is less restrictive than Aumann and Sorin's use of finite memory, because it allows for infinite punishment. A finite memory space is less restrictive than the finite automaton model because it allows for correlation of strategies and randomness.

## 2.6 Large Anonymous Games

Let  $\Gamma$  denote a finite  $n$ -player game. We say that the game  $\Gamma$  is *anonymous* if it satisfies the following conditions. First, the players have identical strategy sets

$$X_i = X \text{ for all } i = 1, 2, \dots, n$$

and, secondly, a player's payoff depends only on his own strategy and the distribution of the strategies chosen by other players

$$u_i : X \times \Delta(X) \rightarrow \mathbf{R} \text{ for all } i = 1, 2, \dots, n.$$

If player  $i$  chooses  $x_i$  and the strategies of players  $j \neq i$  are represented by  $\mu_{-i}$ , then player  $i$ 's payoff is  $u_i(x_i, \mu_{-i})$ . In other words, it does not matter which players choose which strategies. It only matters how many of the players  $j \neq i$  choose a particular strategy. Note that with only  $n$  players, not all the distributions in the set  $\Delta(X)$  correspond to feasible choices for the other players. It is convenient to assume that the payoff function  $u_i$  is defined on all of  $\Delta(X)$  because later we shall want to treat the number of players  $n$  as large and variable.

The corresponding repeated game with limited memory  $\Gamma_L^\infty(s_1)$  will also be anonymous if the transition probability function is anonymous in the sense that it depends only on the distribution of the strategies chosen and not on the profile of strategies:

$$\psi : S \times \Delta(X) \rightarrow \Delta(S).$$

The objective of introducing anonymity is to provide a simple framework in which we can analyze the impact of a single player's actions on the subsequent course of a game as the number of players becomes very large. Note that the memory space  $S$  will be held fixed as the number of players  $n$  increases without bound. If the memory space  $S$  is bounded in an appropriate sense and the number of players is very large, then there is not enough information in the collective memory to make most players pivotal. This is a theme that we shall return to at the end of the chapter.

Eventually, we shall want to apply this analysis to repeated games with limited memory, but first we consider a special case, which turns out to encompass all that we need to know about repeated games with limited memory.

### 2.6.1 The Main Lemma

Suppose that a finite symmetric game  $\Gamma$  is followed by another game that yields a payoff depending only on the state determined by the play of the first game. If the following state is  $s$  then the payoff to player  $i$  is denoted by  $v_i(s)$  and the payoff function is denoted by  $v = (v_1, \dots, v_n)$ . If the play of the first game generates the action profile  $x$  (with distribution  $\mu$ ) then the payoff from the convolution of the two games is given by

$$u_i(x_i, \mu_{-i}) + \delta \int v_i(s) d\psi(\mu, s_1).$$

Of course, player  $i$  does not know the value of  $\mu$  when he chooses  $x_i$  unless all players are choosing pure strategies. To allow for mixed strategies, we assume that  $F_i$  is the probability distribution of  $\mu$  when  $i$  chooses  $x_i$ . Then the payoff to player  $i$  from the continuation game can be written as

$$\pi_i = \int \int v_i(s) d\psi(\mu, s_1) dF_i.$$

We want to show that, under certain conditions, the payoff from the continuation game is independent of player  $i$ 's current choice  $x_i$  when the number of players becomes unboundedly large. To do this, we first have to define a sequence of games and equilibria with increasing numbers of players. An *increasing sequence of games* is defined by

- a finite strategy set  $X$ ;

- a sequence of players  $i = 1, 2, \dots$ ;
- a sequence of individual payoff functions  $u_i : X \times \Delta(X) \rightarrow \mathbf{R}$ , one for each  $i = 1, 2, \dots$ ;
- a symmetric probability transition function  $\psi : \Delta(X) \times S \rightarrow \Delta(S)$ ;
- a sequence of measurable continuation-payoff functions  $v^n : S \rightarrow \mathbf{R}^n$ .

For each  $n$ , the game under consideration consists of the players  $\{1, \dots, n\}$  with the payoff function

$$U_i^n(x_i, \mu_{-i}) = u_i(x_i, \mu_{-i}) + \delta \int v_i^n(s) d\psi(\mu, s_1).$$

We impose the following assumptions on the sequence of games:

- The sequence of functions  $\{v_i^n\}$  is uniformly bounded, that is, for some  $B$  and all  $n$ ,  $|v_i^n| \leq B$ .
- The transition probability function  $\psi$  is assumed to be noisy, in the sense that  $\psi(s, \mu)$  is uniformly absolutely continuous for any  $(s, \mu)$ . To be more precise, if  $\lambda$  is the measure on  $S$  then for some constant  $K$  and any  $(s, \mu) \in S \times \Delta(X)$ ,

$$\psi(s, \mu)(A) \leq K\lambda(A)$$

for any measurable set  $A$  contained in  $S$ .

- $\psi$  is weakly continuous, that is, for any sequence  $\{(s^n, \mu^n)\}$  converging to  $(s^0, \mu^0)$ ,  $\psi(s^n, \mu^n)$  converges weakly to  $\psi(s^0, \mu^0)$ . (See Billingsley (1985), p.335 et seq.).

Fix an arbitrary player  $i$  and consider two possible pure strategies  $x_{i0}$  and  $x_{i1}$ . For each  $n$ , the equilibrium strategies of the other players determine a probability distribution over the action profiles. Let  $F_{ij}^n \in \Delta(X)$  denote the probability distribution over action profiles corresponding to player  $i$ 's choice  $x_{ij}$  for  $j = 0, 1$  respectively. Let  $\pi_{ij}^n$  denote the continuation payoff corresponding to the choice of  $x_{ij}$  in the  $n$ -th game, that is,

$$\pi_{ij}^n = \int \int v_i^n(s) d\psi(\mu, s_1) dF_{ij}^n.$$

We want to characterize the limiting difference  $\pi_{i0}^n - \pi_{i1}^n$  as  $n \rightarrow \infty$ .

To define a limiting equilibrium we must assume that

- $v_i^n \rightarrow v_i^0$  a.e.  $[\lambda]$
- $F_{ij}^n \rightarrow F_i^0$  for  $j = 0, 1$ .

Notice that the second statement implies that in the limit, the distributions are independent of the strategies chosen by player 1. This is obvious, since there is an infinite number of players in the limit. So the only real force of this assumption is to require that the distributions  $F_{ij}^n$  converge to “something”.

With these assumptions, we are finally ready to show that the action of agent  $i$  has a negligible impact on his payoff in the continuation game. Hence, he should choose a best response in the current stage game. The following lemma formalizes this claim.

**Lemma 8** *Under the maintained assumptions,  $(\pi_{i0}^n - \pi_{i1}^n) \rightarrow 0$  as  $n \rightarrow \infty$ .*

**Proof.** Now

$$\begin{aligned} & \int \int v_i^n(s) \psi(ds|s_1, \mu) F_{ij}^n(d\mu) - \int \int v_i^0(s) \psi(ds|s_1, \mu) F_i^0(d\mu) \\ = & \int \int (v_i^n(s) - v_i^0(s)) \psi(ds|s_1, \mu) F_{ij}^n(d\mu) - \\ & \int \int v_i^0(s) \psi(ds|s_1, \mu) (F_{ij}^n(d\mu) - F_i^0(d\mu)) \end{aligned}$$

so that

$$\begin{aligned} & \left| \int \int v_i^n(s) \psi(ds|s_1, \mu) F_{ij}^n(d\mu) - \int \int v_i^0(s) \psi(ds|s_1, \mu) F_i^0(d\mu) \right| \\ \leq & \int \int |v_i^n(s) - v_i^0(s)| \psi(ds|s_1, \mu) F_{ij}^n(d\mu) + \\ & \left| \int \int v_i^0(s) \psi(ds|s_1, \mu) (F_{ij}^n(d\mu) - F_i^0(d\mu)) \right|. \end{aligned}$$

By Egoroff’s theorem (Royden (1988), p.73, Exercise 30), we know that  $v_i^n$  converges almost uniformly to  $v_i^0$ , that is, for any  $\varepsilon > 0$  there exists a set  $A \subset S$  such that  $\lambda(A) \leq \varepsilon$  and  $v_i^n$  converges uniformly to  $v_i^0$  on  $S \setminus A$ . Choose  $N_1$  so that

$$|v_i^n(s) - v_i^0(s)| \leq \varepsilon$$

for all  $n \geq N_1$  and  $s \in S \setminus A$ . For  $s \in A$  we know that  $|v_i^n(s) - v_i^0(s)| \leq B$  for all  $n$ , so

$$\int \int |v_i^n(s) - v_i^0(s)| \psi(ds|s_1, \mu) F_{ij}^n(d\mu) \leq \varepsilon + BK\varepsilon$$

for  $n \geq N_1$ . By Lusin's Theorem (Halmos (1974), p.242),  $v_i^0$  is almost continuous, that is, there exists a compact set  $C \subset S$  such that  $\lambda(S \setminus C) \leq \varepsilon$  and  $v_i^0$  is continuous on  $C$ . Therefore, from the definition of weak convergence and from the fact that  $F_{ij}^n$  converges weakly to  $F_i^0$ , there exists a number  $N_2$  such that

$$\begin{aligned} & \left| \int \int v_i^0(s) \psi(ds|s_1, \mu) (F_{ij}^n(d\mu) - F_i^0(d\mu)) \right| \\ & \leq \left| \int \int_C v_i^0(s) \psi(ds|s_1, \mu) (F_{ij}^n(d\mu) - F_i^0(d\mu)) \right| + \\ & \quad \left| \int \int_{S \setminus C} v_i^0(s) \psi(ds|s_1, \mu) (F_{ij}^n(d\mu) - F_i^0(d\mu)) \right| \\ & \leq \varepsilon + BK\varepsilon \end{aligned}$$

for all  $n \geq N_2$ . Putting the two inequalities together, we have

$$\left| \int \int v_i^n(s) \psi(ds|s_1, \mu) F_{ij}^n(d\mu) - \int \int v_i^0(s) \psi(ds|s_1, \mu) F_i^0(d\mu) \right| \leq 2(\varepsilon + BK\varepsilon)$$

for  $n \geq N = \max\{N_1, N_2\}$ . Since  $\varepsilon$  is arbitrary, it follows that  $\pi_{ij}^n \rightarrow \pi_i^0 \equiv \int \int v_i^0(s) \psi(ds|s_1, \mu) F_i^0(d\mu)$  as  $n \rightarrow \infty$  for both  $j = 0, 1$ . ■

The proof of Lemma 8 is intended to be heuristic rather than efficient. In particular, it is interesting to see how two of the three principles of functional analysis, Egoroff's theorem and Lusin's theorem, come into play in the proof. Needless to say, the whole argument could be replaced as a single continuity assumption, but that would have been very unilluminating.

Sabourian (1990) summarizes his central technical proof as follows:

“The proof is based on

(a) showing that in any finite-player repeated game, the maximum gain a player can make, in any period, by deviating from any Nash equilibrium path of the repeated game, is bounded above by a product term consisting of a fixed number and the (metric) distance (when total variation norm is used) between the distribution of the random outcome when a deviation takes place, and that when the deviation does not take place;

(b) showing that anonymity and continuity imply that the above distance between the two distributions approaches zero as the number of players becomes large;

(c) constructing a sequence of finite player games that converge to a continuum, using Skorohod’s embedding theorem (often used in core theory).” (Sabourian (1990; pp.95-96)).

Lemma 8 corresponds to parts (a) and (b) of this proof strategy. Without going further to carry out part (c), we cannot define a limit game or say in what exact sense the Limit Principle holds in this framework. However, it is intuitively clear that “in the limit” each player will choose an action, contingent on any memory state  $s$ , that is a best response in the current stage game. In this sense, what we should expect to observe is a profile of actions that forms a Nash equilibrium of the stage game at every date  $t$  and in every memory state  $s$ . This implies several different features of the equilibrium (in the limit) that are relevant to our enquiry:

- In choosing his action at any date  $t$  and in any memory state  $s$ , an agent does not take into account the effect that his action may have on the future play of the game. This is because, asymptotically, his action has a negligible effect on the next memory state and hence a negligible effect on the future play of the game.
- Although the equilibrium does not have the Markov property (the agents’ strategies depend on the entire history of public signals as well, possibly, as their own actions), it does share the crucial feature of Markov equilibria, namely that punishment strategies based on an individual player’s deviation are ruled out. Markov strategies rule out punishment strategies by preventing any dependence on memory; here we achieve the same end by restricting the dependence to aggregate signals and making sure that this dependence is continuous as an individual player becomes vanishingly small.
- This does not mean, however, that the memory states do not matter. The memory states can serve as a correlation device and random evolution of the memory states, independently of the actions of the agents, can generate “sunspot phenomena”. What we observe over time is equivalent to a correlated equilibrium in the sense of Aumann (1974).

What might be called the Continuity Principle, the property that a single agent’s action has a vanishingly small effect on the play of the game when the number of agents is unboundedly large, has been justified by introducing

(a) a small amount of noise and (b) a sufficiently small memory space. These are the kind of assumptions that we need for to justify the conditions used in Chapter 1.

## 2.7 Non-Anonymous Games

While anonymity seems a natural assumption in the context of many models of markets, it is also restrictive. First, ‘anonymity’ is one of the properties that we would like to obtain endogenously from a study of the structure of the market game, rather than imposing it by assumption on the market game. Second, anonymity also reduces the dimensionality of the memory space by aggregating the actions of the agents into a distribution that determines the evolution of the state of the game. Expanding the dimensionality of the state space may change the results. A third restriction of the preceding analysis does not follow from anonymity per se, but is associated with it: in order to derive the limit principle we have to assume that the transition probability is *uniformly* absolutely continuous. What this means in more accessible terms is that the noise introduced into memory becomes large relative to the size of an individual agent’s impact on the economy. While this does not immediately imply that the individual agent’s action is not pivotal, it might be suspected of having something to do with it. In particular, it would be interesting to see whether smaller amounts of uncertainty would do the job.

In this section, another approach is adopted to test the robustness of the results based on anonymous games. The essential ideas are drawn from a related problem in the literature on mechanism design and the provision of public goods.

### 2.7.1 Continuity in the Provision of Public Goods

Suppose there are  $n$  individuals  $i = 1, \dots, n$  in a community which has to decide whether to undertake a project to provide a public good. The project is indivisible, so without loss of generality we can assume that the amount of the public good is either 0 or 1. Preferences are linear in money. Then we can normalize the utility function so that each individual gets a payoff of  $v - p$  if the good is provided, where  $v$  is the value of one unit of the good and  $p$  is the individual’s contribution to the provision of the public good. If the good is not provided, each individual receives a payoff of 0.

There are two types of individuals, depending on the value assigned to the public good. Some get utility  $v_H$  from the good and some get utility  $v_L$ , where  $0 < v_L < v_H$ . Individuals are ex ante identical and have independently distributed types, with probability  $0 < \alpha < 1$  that the value of the public good is high. The total cost of the good is  $nc$ , where

$$v_L < c < (1 - \alpha)v_L + \alpha v_H.$$

In other words, the average cost of the public good is greater than the low valuation, so an agent with a low valuation will not want the good if he has to pay the average cost; but the expected value of the good to a typical agent is greater than the average cost, so ex ante it is efficient to provide the good.

We assume that individuals can opt out of the project: if they are asked to contribute more than the good is worth to them, they can refuse to participate.

Let  $v_i$  denote the valuation of agent  $i$ , let  $p_i$  denote the contribution of agent  $i$ , and let  $x$  denote the quantity of the public good provided. A mechanism for providing the good is a function that maps a profile of messages  $(v_1, \dots, v_n)$  into a decision about the provision of the good  $x = 0, 1$  and a profile of contributions  $(p_1, \dots, p_n)$  such that  $\sum_{i=1}^n p_i \geq xnc$ . The individual rationality constraint requires that  $v_i - p_i \geq 0$  for all  $i$  where  $v_i$  is the actual valuation of agent  $i$ . By the Revelation Principle (see, e.g., Myerson (1991), p.260) we can restrict attention to truth-telling equilibria. The problem faced by the mechanism designer is that, given the individual rationality constraint, individuals are likely to want to misrepresent their preferences, claiming to have a low valuation in order to be assigned a low contribution. Why would anyone would want to participate in the provision of the public good if by opting out they can avoid paying anything? The answer is that the individual may regard himself as pivotal in determining whether the public good is provided or not. We can see this point through a simple example.

Suppose there are only two individuals,  $i = 1, 2$ . If both have low valuations it is not worth providing the good; if both have high valuations it is worth providing the good; and if one has a high valuation and one has a low valuation it is worth providing the good if

$$2c < v_L + v_H,$$

which I assume to hold in what follows. Consider the following mechanism:

- if both announce low valuations, the good is not provided and no contributions are paid;
- if both announce high valuations, the good is provided and both pay  $c < v_H$ ;
- if one announces the high valuation and one announces the low valuation, the good is provided, the agent with the low valuation pays  $v_L$  and the agent with the high valuation pays  $2c - v_L$ .

By making the low-valuation agent pay the maximum consistent with his individual rationality constraint, we relax the incentive constraint on truth-telling as much as possible. This mechanism is clearly feasible, individually rational, and efficient if agents tell the truth. Now consider the incentives for the agents to tell the truth.

If an agent has the low valuation, it cannot be in his interest to report the high valuation. At best this will raise his contribution without changing the provision of the public good; at worst it will raise his contribution and cause inefficient provision of the public good. So suppose that the agent has the high valuation and considers reporting the low valuation. If the other agent announces the high valuation, reporting the low valuation does not change the provision but reduces the contribution. If the other agent announces the low valuation, announcing the low valuation ensures that the good is not provided, and this reduces the agent's payoff. So there is a tradeoff between reducing one's contribution and reducing the probability of providing the good. The agent knows that in some states he is pivotal in determining the provision of the public good and this disciplines him against mis-reporting his valuation.

The payoff from mis-reporting when the true valuation is  $v_H$  is

$$\alpha(v_H - v_L)$$

and the payoff from truth-telling is

$$v_H - (1 - \alpha)(2c - v_L) - \alpha c.$$

The incentive constraint then is

$$\alpha(v_H - v_L) \leq v_H - (1 - \alpha)(2c - v_L) - \alpha c$$

or

$$(1 - \alpha)v_H + v_L \geq (2 - \alpha)c.$$

This constraint is satisfied if the surplus is high and the value of  $\alpha$  is not too high. For example, if  $\alpha = 1$  then the constraint is clearly incompatible with the assumption that  $v_L < c$ .

What is true for small numbers of agents may not be true for large numbers, however. As the number of players gets larger, providing the public good is efficient as long as the fraction of agents with a high valuation is at least  $\alpha$ . A single agent's signal will not change the distribution of signals reported very much and so the chance of the agent being pivotal seems intuitively to become small. When the number of agents is very large, the law of large numbers ensures that the fraction of agents with the high valuation will be very close to  $\alpha$  with high probability, so the good will be produced with probability close to one if agents tell the truth. The temptation for an agent with a high valuation to mis-report his signal becomes very strong. In fact, it is easy to see that the public good will not be provided with probability one as efficiency requires. Beyond this obvious statement, however, it is hard to say exactly what will happen. As long as an agent thinks that he is pivotal, he may have an incentive to tell the truth. Furthermore, for any large but finite number  $n$ , some agents can be pivotal. What is not obvious is how many agents can be made pivotal and how much of the public good can be provided.

This problem was studied by Mailath and Postlewaite (1990) who showed that as the number  $n$  becomes unboundedly large, the probability that the public good is provided converges to zero for any individually rational and incentive compatible mechanism. The main step in proving this result is a lemma that shows that, as  $n$  grows unboundedly large, almost all agents become non-pivotal. A particularly elegant proof of this principle is contained in an appendix to Mailath and Postlewaite (1990) and is attributed by them to a referee of the paper. In what follows, I am going to adapt the referee's argument to a game-theoretic environment to show that even in games that are not symmetric (anonymous), the Continuity Principle holds for almost all players as the number of players becomes unboundedly large.

### 2.7.2 The Referee's Lemma

Suppose that there is a countable sequence of players  $i = 1, 2, \dots$  and that each player has a finite action set  $X_i$ . For any  $n$  we define the game  $\Gamma^n$  to consist of the first  $n$  players  $i = 1, \dots, n$ . Without the assumption of anonymity, the payoff function of a given player changes as the number of players varies. Let  $u_i^n$  be a real valued function defined on  $X^n \equiv X_1 \times \dots \times X_n$  and let  $u^n \equiv (u_1^n, \dots, u_n^n)$  be the payoff function for  $\Gamma^n$ .

As before, the payoff from the continuation game is represented by the payoff functions  $v_i^n : S \rightarrow \mathbf{R}$  for every  $i \leq n$  and  $n = 1, 2, \dots$

The evolution of the state is determined by the transition probability function  $\psi^n : X^n \rightarrow \Delta(S)$ , where reference to the initial state is suppressed because it is held fixed in what follows. We assume that  $S$  is (contained in) a Banach space. (See Royden (1988), p.217).

For each  $n$  the equilibrium strategies have the distribution  $F^n \in \Delta(X^n)$ .

Now consider a sequence of games and the corresponding equilibria satisfying the following conditions:

- $v_i^n$  is a measurable function and for each  $i$ , the pointwise limit  $\lim_{n \rightarrow \infty} v_i^n = v_i^0$  is a measurable function;
- for each  $n$  the function  $\psi^n$  is continuous when  $\Delta(S)$  is endowed with the topology of weak convergence; furthermore,  $\lim_{n \rightarrow \infty} \psi^n = \psi^0$  in the sense that for any sequence  $\{x^n\}$  converging to  $x^0$ ,  $\psi^n(x^n)$  converges weakly to  $\psi^0(x^0)$ ;
- there exists a distribution  $F^0$  on  $X^0 = X_1 \times X_2 \times \dots$  such that the marginal distributions of  $F^n$  converge to the marginal distributions of  $F^0$ .

Our interest is first with the limiting “equilibrium” represented by  $F^0$ . The independence of the strategies used by the players implies that for each  $n$

$$F^n = \times_{i=1}^n F_i^n$$

so in the limit we have

$$F^0 = \times_{i=1}^{\infty} F_i^0,$$

where we can interpret the marginal distribution  $F_i^0$  as the distribution of player  $i$ 's mixed strategy. Let  $(\Omega, \mathcal{B}(\Omega), P)$  be a probability measure space

and for each  $i$  let  $x_i$  be a random variable defined on  $(\Omega, \mathcal{B}(\Omega), P)$  with values in  $X_i$  and with marginal distribution  $F_i^0$ . Similarly, let  $y$  be the random variable on  $(\Omega, \mathcal{B}(\Omega), P)$  with values in  $S$  and having a distribution given by  $\psi^0 \circ F^0$ . Then define the random variable  $y_i$  to be the expectation of the state  $y$  conditional on player  $i$ 's mixed strategy  $x_i$ :

$$y_i = E[y|x_i]$$

and normalize the random variables  $y_i$  by putting

$$\bar{y}_i = \frac{y_i}{E[(y_i)^2]}$$

so that  $E[(\bar{y}_i)^2] = 1$  and  $E[y_i] = E[y]$ . Without loss of generality we can normalize  $y$  so that  $E[y] = 0$  in what follows so that  $E[y_i] = 0$  for each  $i$ .

Now, it is important for what we are going to do that each of the random variables  $y_i$  be distinct. For this to be true, the random variables must be non-degenerate. There are two ways that this may happen. First, if  $x_i$  is degenerate then, of course,  $y_i$  will be degenerate too. So we must assume that  $x_i$  is not degenerate. However, even if  $x_i$  is non-degenerate it may still be the case that  $y_i = 0$  if  $x_i$  does not influence  $y$ . However, since we want to prove that  $x_i$  has no influence on  $y$ , there is no loss of generality in assuming that this case does not arise. Equivalently, disregard the agents for whom this is true. So the critical assumption is that  $x_i$  is non-degenerate, i.e., not a pure strategy.

It is easy to see why pure strategies need to be ruled out. Suppose that

$$X_i = \{0, 1\}, \forall i,$$

$$S = \{0, 1\},$$

$$\psi^0(x) = \max_i \{x_i\}, \forall x \in X^\infty$$

and

$$v_i^0(s) = -sM$$

for some large number  $M$ . Then it will be easy to support the choice of  $x_i^* = 0$  for all  $i$  by choosing  $M$  sufficiently large. Any deviation from  $x_i^* = 0$  by a single player will subject the player to a large punishment that offsets any short run gain. Despite the fact that there is an unboundedly large number of players, each player regards himself as pivotal and it is this perception that supports the choice of  $x_i^* = 0$  as an equilibrium.

Note that even a small amount of uncertainty will upset this equilibrium. For example, suppose that we appeal to the “trembling hand” argument of Selten (1975) to test the robustness of this equilibrium. We introduce a small “tremble” by requiring each agent to place a probability weight of at least  $\varepsilon > 0$  on each pure strategy. However small  $\varepsilon$  is, the number of agents choosing  $x_i = 1$  will be infinite with probability one. To make each player pivotal in this way requires each player to be standing on the edge of a precipice so that a small amount of uncertainty will push some of them over. Likewise, if players choose a mixed strategy, then  $\max\{x_i\} = 1$  with probability one so again the action of an individual player has no effect.

Note also that in a anonymous game it is only the distribution of the players’ actions that counts and the effect of the trembling hand disappears in the limit. The fraction of agents choosing each strategy will be almost surely constant by the strong law of large numbers. So the uncertainty introduced by the trembling hand would not have been sufficient to produce the desired result in the preceding section. However, it is sufficient in what follows. I assume therefore that for some  $\varepsilon > 0$  and all  $i$  the probability placed on each strategy by  $F_i^0$  is at least  $\varepsilon$ :

$$P[x_i = \xi] \geq \varepsilon, \forall \xi \in X_i.$$

Since the strategy of player  $i$  puts positive probability on each action (pure strategy), there is a well defined expectation of the state conditional on any choice of action by player  $i$ . This allows us to see the effect that any choice of action by player  $i$  would have on the continuation payoff and determine whether player  $i$  is pivotal or not.

Since the random variables  $x_i$  are independent, the random variables  $y_i$  are also. If we assume that  $y_i$  is non-degenerate, i.e., not identically 0, then it follows that all the elements of  $\{\bar{y}_i\}$  are distinct.

Consider the random variables  $\{\bar{y}_i\}$  and  $y$  as elements of  $L^2((\Omega, \mathcal{B}(\Omega), P), S)$ , that is, the linear space of square integrable functions from  $(\Omega, \mathcal{B}(\Omega), P)$  to  $S$ . (The mathematical basis for the following is found in Royden (1998), p.245 et seq.). With the inner product defined by

$$x \cdot y = \int xy dP,$$

$L^2$  is a Hilbert space and  $\{\bar{y}_i\}$  is an orthonormal system, that is, the elements are orthogonal

$$\bar{y}_i \cdot \bar{y}_j = 0 \text{ for all } i \neq j;$$

they are distinct

$$y_i \neq y_j \text{ for all } i \neq j;$$

and they have unit norms

$$\|\bar{y}_i\| \equiv \int (\bar{y}_i)^2 dP = 1 \text{ for all } i.$$

Define the Fourier coefficients  $\{a_i\}$  by putting

$$a_i = E[y\bar{y}_i]$$

for each  $i$ . Then Bessell's Inequality (Royden (1980), p.246) tells us that

$$\sum_{i=1}^{\infty} a_i^2 \leq \|y\|^2 = E[y^2] < \infty.$$

In other words,  $a_i \rightarrow 0$  as  $i \rightarrow \infty$ . This means that for any  $\varepsilon > 0$ ,  $|a_i| < \varepsilon$  for all except a finite number of players  $i$ . Now  $a_i$  measures the impact of  $x_i$  on the memory state  $y$ . More precisely,  $a_i$  measures the correlation between the state  $y$  and the conditional expectation of the state  $y_i$ . When  $a_i$  is small, the action chosen by agent  $i$  has little value in predicting the outcome  $y$ .

A more general treatment of these problems has been provided by Lehrer and Neeman (1998) and Al-Najjar and Smorodinsky (1998a) in the context of mechanism design problems. Al-Najjar and Smorodinsky (1998b) have extended the earlier results of Green (1980) and Sabourian (1990) to allow for games in which the aggregate outcome is not necessarily an anonymous function of players' actions and where players' strategies may depend non-anonymously on signals of other players' actions.

### 2.7.3 Pivotal Players in Games

Speaking loosely, the Referee's Theorem tells us that all but a finite number of agents will have little effect on the outcome of the game, so almost every player will be choosing an action that maximizes his payoff in the stage game, independently of what happens in the continuation game. This is a very powerful and general result, but it is important to see what it does and does not mean. The following example will be useful.

Consider the following two stage game. There is a large but finite number of sellers  $S$  and a large but finite number of buyers  $B > S$ . Every seller has

one unit of an indivisible good, whose value to the seller is zero, and every buyer wants to purchase one unit of the good, whose value to the buyer is one. All agents have quasi-linear utility functions. In the first stage of the game, the buyers and sellers are randomly matched and the seller makes a proposal to the buyer (unmatched buyers must remain passive). The seller can offer one of two prices  $0 < p_L < p_H < 1$ . The buyer accepts or rejects the offer. If the offer is accepted, the trade is carried out at the proposed price; the buyer gets a payoff of  $1 - p$  and the seller gets a payoff of  $p$ . The agents who have traded leave the game and take no further part. The rest of the players, whether they were matched or not, proceed to the second stage of the game. At the second stage, the matching of the remaining buyers and sellers and the prices at which the goods are exchanged between each pair are determined by an exogenous mechanism. Again, if a buyer and seller trade the good at the price  $p$ , the buyer's payoff is  $1 - p$  and the seller's payoff is  $p$ . Agents who do not trade get 0.

The mechanism that comes into play at the second stage is, of course, intended to represent the continuation of the game without forcing us to analyze a (potentially) infinite-horizon game. By choosing an appropriate mechanism at the second stage, we can easily manipulate the agents' payoffs in any way we like. Suppose, for example, that we wanted the goods to trade at the price  $p_L$ . If a seller demands a price  $p_H$  then the buyer rejects and next period gets the good at a price of zero. If the seller demands  $p_L$  and the buyer rejects, then the buyer gets the good in the next period for a price of one. In the same way, we could support any mixture of high and low prices by appropriately punishing individual buyers and sellers for their deviations at the first stage. It does not matter how many agents there are in the game, each agent perceives that his actions are going to have a marked effect on the future of the game. In what ways does this example violate the conditions of the Referee's Theorem?

First of all, we have only considered pure strategies, whereas the Referee's Theorem requires us to have completely mixed strategies. The example is easily adapted to allow for mixtures by introducing a trembling hand. Suppose that with probability  $\varepsilon$  the price offered is the opposite of the one chosen by the seller. For  $\varepsilon > 0$  small enough, this will have no effect on the equilibrium described above. If the trembling hand forces the seller to offer the wrong price, he will be punished at the second date, but this will not alter his best response because the probability of punishment is small and exogenous. Similarly, if there is a small probability that the buyer will accept

when he should reject or reject when he should accept, that will not change the seller's best response if the probability of the tremble is sufficiently small. So the absence of mixed strategies does not by itself explain the existence of multiple equilibria here.

The second point to note is that the space of outcomes from the first stage of the game includes, for each pair consisting of a buyer and a seller, the price offered and the response. The dimension of the memory space required to describe this outcome is proportional to the number of agents. As the number of agents increases, the dimension of the memory space increases. With more buyers and sellers, there are more goods to trade, more prices to choose, and more payoffs to determine. In the framework used to describe the Referee's Theorem, there is a fixed memory space. As the number of agents increases, the dimension of the space of action profiles increases, but the range of the function that maps action profiles into memory states or outcomes remains the same. So a larger and larger action space is being mapped into a fixed memory space. To do this in a reasonably orderly way, one has to ignore most of the information in the larger space, that is, most of the components of the action profile do not matter very much in determining the outcome.

In a public good provision problem, the outcome that agents care about, the provision of the public good, does not change as the number of agents increases. So it is natural to work with a fixed outcome space. This is not true in a market with private goods, where the dimension of the allocation increases with the number of agents. In the bargaining game above, two things are happening as the number of buyers and sellers increases. First, the dimension of the memory space increases, allowing the actions of more agents to be recorded in the memory state. Secondly, the part of the memory state that a single agent cares about is getting smaller relative to the outcome or memory state. A typical buyer, for example, does not care about the entire allocation; he only cares about getting one unit of the good for himself and the price he has to pay for it. An individual agent's action has little effect on the outcome of the first stage, in the sense that his trade and the corresponding price are a negligible part of the resulting description in the memory state; but the agent's action has a big effect on his payoff in the continuation game. This is why the seller does not maximize his short-term payoff in the example.

To rule out this kind of equilibrium, in which a seller's first-stage action has a large effect on his second-stage payoff, it is not enough to hold the memory space constant as the number of agents increases. One could always

achieve the same effect by assuming the memory space  $S$  is so large that it accommodates a complete description of the first stage outcome for any number of players. So in addition one needs to assume that  $S$  is not “too large” and that in turn requires some structure in terms of the continuity properties of the payoffs  $v_i^0$  and the topology on  $S$ . One needs both the assumption that a single agent’s action has a small effect on  $s$  and the assumption that a small change in  $s$  has a large effect on  $v_i^0$ . The second assumption is violated in the example, because the action of the seller has a small effect on the state but a large effect on  $v_i^0$ . To rule this out, we could assume that the payoff functions  $\{v_i^0\}$  are in some sense equi-continuous. This again is a form of anonymity, a restriction on the extent to which individual actions can be remembered and used to condition the future payoffs of individual agents. In the end we are assuming that an agent’s future payoff is dependent on an aggregate state which the agent’s current action can only effect marginally.

What this example makes clear is that, even in a non-anonymous game, some form of the anonymity assumption is required.