1 Introduction

The previous chapter concentrated on finding solutions to Markov games that are characterized by the property that previous choices made in these dynamic games affect current choices only indirectly, if at all, through the current state of play. This still leaves open the question as to whether there are any other solutions. Part of the answer depends on whether there is more than one Markov solution, the other part on whether the game ends after a finite number of moves or not. To give a more definitive answer we turn to the special case of stage games.

In stage games the probability transitions linking the stages are not affected by the choices players make. It follows that players cannot affect the trade off between current payoffs and future payoffs except through their strategic interactions. It is natural to ask if all the solutions to the game can be found by treating each stage as an separate game and piecing together the solutions of the individual stages. Section 3 answers this question. If the probability transition is independent of all the choices made, if the game is finite (has a finite number of decision nodes), and if every stage has a unique solution, then the unique solution to the game is found by piecing together the unique solutions to the stage games partitioning it. In these cases we argue there is no role for strategic investment

Scope for strategic investment arises when either of the two latter conditions are violated, and the remainder of the chapter is concerned with what happens when uniqueness of the kernel games do not have a unique solutions, or when the game has an infinite horizion. Our discussion is focuses on repeated games, that is a game subjects played several times with each other without switching roles, but the main results carry over to stage games without much revision.

If the choices players make affect the probability transitions of reaching subsequent stages, then the incentives for playing within a stage do not necessarily align with the player's incentives within the game, because there may be a trade-off between the current accruals from completing the stage and the benefits arising from the stages which follow. We have already studied this phenomenon in teh previous chapter in our discussion of Markov solutions.

When the probability transitions depend on player choices, it is easy to understand why ignoring everything but current accruals might be very misleading. More surprising is the fact that even when all the transition probabilities are exogenous, concentrating on the set of solutions generated by the stage solutions provides a very incomplete picture of the solution set for the whole game, the topic of discussion in Sections 5 and 6.

These two sections also provide insight into the nature of coordination, leadership and reputation. We explain why these concepts are hard to meaningfully define in games with unique solutions, and they are often mentioned in the context of repeated interactions. For example, coordination is not spontaneous, typically requiring the same players to repeatedly interact with each other in order to converge to stable outcomes, a remark we amplify. Firms, both small and large, develop reputations for product quality and after sales service through dealings with successive customers. Similarly, retail and service chains and franchises develop reputations for consistency in their product offerings across different outlets. Politicians cultivate their aspirations for leadership through their repeated personal interactions within a public forum. In Sections 5 and 6 leadership is defined as a quality that facilitates coordination to reach one of several solutions , that is a self enforcing agreement in a game that supports multiple solutions. Similarly reputations are established and maintained when players reach an implicit understanding about which solution is being played. We treat the tacit selection of a pure strategy solution achieved through coordination as a self enforcing agreement, a topic we will revisit in Chapter 14 on contracts.

2 Unique Solutions to Stage Games

Much of the intuition and some of the results on repeated games also apply to the whole class of stage games. Thus our discussion of stage games in this chapter will focus on repeated games. It is natural to ask whether the predicted outcomes of the repeated game can be found by simply repeating the predictions of a typical round. Yet because the strategy space for the repeated game is more than the product of the number of strategies in one round and the number of rounds, it is not obvious why this should be the case. In the sections which follow we seek to resolve this issue. First we focus on games which are repeated a finite number of times, then on infinite horizon games where payoffs are discounted over time.

We begin our analysis of the solution to games with stages

2.1 Textiles

We first consider repeated games in which the kernel has a unique equilibrium. Figure 10.1 depicts the strategic form of a textile industry, where two producers must decide between using local unionized labor force in production versus foreign child labor. If either one of the firms chooses the to use local labor, the other one will be heavily penalized from tariffs supported by organized labor ostensibly to punish the exploitation of impoverished children by multinationals. Furthermore the firm using unionized local labor incurs substantial penalties by union action from switching its textile mills to overseas. We suppose the kernel game is repeated 3 times.

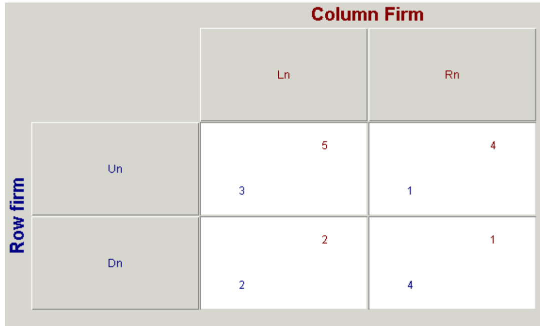

Figure 12.5
Textile production

The repeated game can be solved using the principle of subgame perfection developed in Chapter 8. The smallest subgame in this repeated game is the simultaneous move games which begins at the third period. This subgame corresponds to the kernel game itself. It is evident from Figure 12.5 that the unionized firm has a dominant strategy to use local labor. Applying the rule of iterated dominance, the non-unionized firm uses local labor as well. This strategy profile constitutes the unique equilibrium yielding a payoff vector of $(3,5)$ to the respective firms.

Consider now the strategic form of the reduced subgame beginning in the second period. To compute the payoffs in each cell we add a payoff of 3 to the payoffs that the non-unionized firm receives, and a payoff of 5 to the unionized firm. Here the critical point to notice is that because we are adding a constant number to each of the payoffs each firm receives one of then inequities that determine equilibrium in the kernel game are affected. Folding back, the strategic form of the reduced game starting at period 2 is given in Figure 12.6.

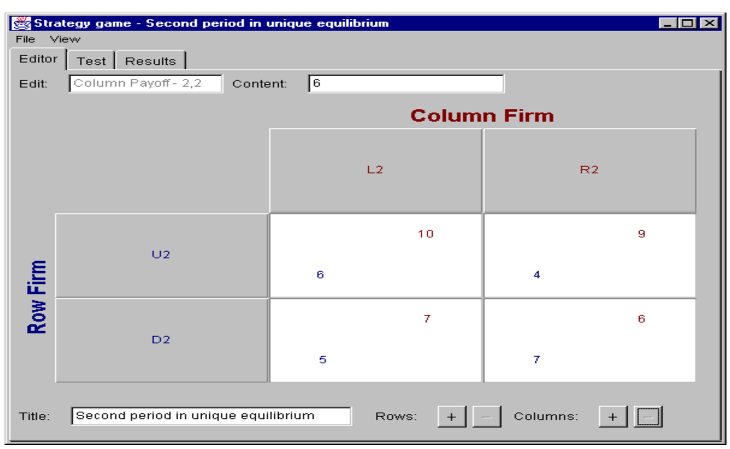

Figure 12.6
Reduced subgame for textile industry

Thus drawing upon the local labor force is also a dominant strategy for the unionized firm in the second period, and the principle of iterated dominance applies equally to the nonunionzed firm. Once again the unique equilibrium for the reduced game is for neither firm to employ foreign child labor, and the payoffs to both firms are $(6,10)$.

The strategic form for the reduced game starting at the beginning of the first period are derived in a similar manner to the strategic form for the reduced game starting in the second period. This time we add the vector of $(6,10)$ to the payoffs in the cells of the kernel game's strategic form. See Figure 10.3. The solution is derived in exactly the same way. We conclude that both firms employ local labor for all three periods of the game.

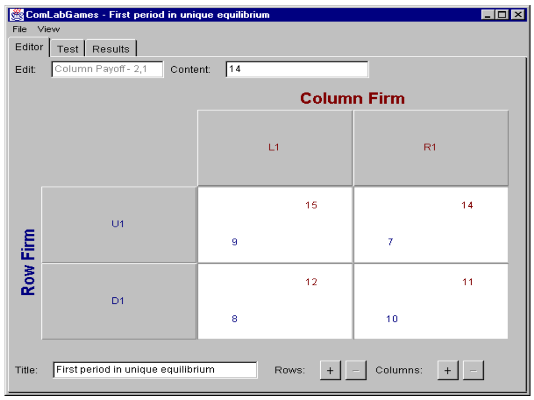

Figure 12.7
The reduced game in strategic form

2.2 A general result

The fact that the solution to the kernel of the textile game is unique explains why such a tight characterization of the solution to the repeated game is obtained. As we fold back the solution of a subgame to form the reduced game that precedes it, the same value is added to each cell of the kernel game for any given player. As we have seen this value might differ across players, but it does not depend on which move a given player chooses in the preceding reduced game. This feature is also present in stage games that consist of kernels with unique equilibrium. Consequently a finite stage game with a unique equilibrium supporting each kernel game, supports a unique equilibrium, formed from the sequence of the kernel equilibrium. We state this result as a theorem .

Theorem *Suppose every kernel game of a finite stage game has a unique solution. Then the stage game also has a unique solution, which is the sequence of solution to the kernel games.*

These examples and results show that neither reputation nor leadership count when all the kernel games in a finite horizon stage game have a unique equilibrium.

Reputations and leadership can only arise when at least one of the following three conditions is present:

1. The probability transitions are choice specific

2. There are multiple solutions to at least one of the kernel games.

3. The kernel games are repeated indefinitely.

The rest of this chapter now explores these three situations in detail.

3 Multiple Equilibrium in Finite Horizon Repeated Games

Having investigated finite horizon stage games in which every kernel game has a unique equilibrium, there are two directions to take. One direction is to explore relaxations of the requirement that every kernel game supports a unique equilibrium. The other is to extend the analysis to infinite horizon stage games. This section pursues the first direction; in the next section we take the second direction. In both sections we concentrate on repeated games, although the basic intuition carries over to the larger class of stage games. In the previous chapters we already encountered games with multiple solutions, but have deferred until now an extended discussion about them. We begin this section with a word of caution, that why players might settle at one solution versus another, and how quickly they converge, are currently topics of debate for researchers, who might ultimately lead us to doubt whether solution refinements beyond backwards induction and dominance principles are very compelling if they are not unique.

Studying coordination games provides a framework for exploring about these issues, and these are the first items on our agnnda. In coordination games all the solutions have identical payoffs, or more generally all the outcomes are ranked by players the same way, so they provide an ideal medium for studying how easily players can alight on a focal point of mutual gain.

3.1 Coordination

Coordination games are marked by the absence of conflict between players. Their objectives are fully aligned. This does not imply they receive the same payoffs, only that players rank all the solution outcomes the same way, from the best to the worst, and therefore could have the same utility function. Coordination games are frustrating to play if there are multiple solutions to the game that tie for the top rank. Unless players agree in advance upon a specific solution, then they all receive a lower payoff than they would attain if there is no opportunity to reach a prior accord.

Consider, for example, the following game called Coffee Break. Suppose Romeo and Juliet have no means of communicating their intense feelings for each other to arrange a time and place to meet. In the following game if they both take coffee at the same time and place, then each has an excuse to engage in small talk. Otherwise no meeting takes place. The strategic form of the game is illustrated below.

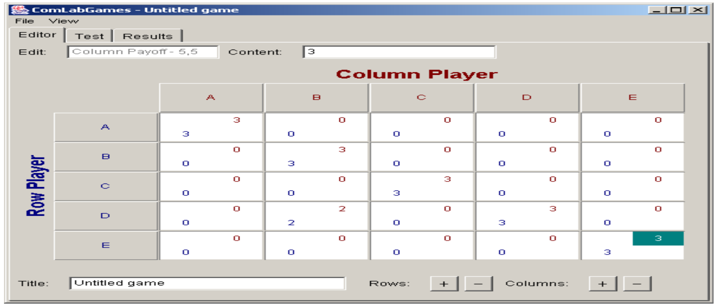

Figure 12.12
A coordination game

There are ten pure strategy equilibrium (and many more mixed strategy equilibrium, all of which achieve lower payoffs). Furthermore every choice is part of exactly one pure strategy equilibrium. When will a spontaneous meeting occur? If each player initially chooses a time randomly, then the probability of meeting each other is one tenth.

One reason why players do not behave according to the predictions of game theory is that the solution concept requires coordination. For example consider Figure 11.1, which displays a simultaneous move game where there is no conflict between players about the objective but ambiguity about how to achieve it. Applying the analysis from Chapter 8, there are two pure strategy Nash equilibrium achieving a unit payoff, and a third mixed strategy equilibrium where players mix with equal probability, yielding an expected value of one half each. If the players collectively pick top left or bottom right they are both rewarded with one unit, but otherwise neither receives anything. Unless prior communications between players permits them to reach an understanding about how the other one will move, it is not obvious why either player would make a particular choice. The principles we have developed nothing to say which equilibrium will be selected and whether repetition will lead to the same outcome. Are we to assume that no systematic patterns will emerge in the data?

The simple example shows that imperfect information games supporting pure strategy equilibria containing inner loops that are not dominance solvable pose a further challenge to players even if they are not complicated to solve, because they require more coordination to reach those equilibrium. In this case there are no general rules for guiding the coordination, and so the knowledge is sometimes specific to the team or players within the game, within a class, a community or perhaps a culture. By definition coordination is not spontaneous, typically requiring the same sets of players to repeatedly interact with each other, although not necessarily in the same roles. Studying repeated games gives us the opportunity to investigate the phenomenon of coordination.

An alternative hypothesis is that the mixed equilibrium strategy profile will be played every round until a positive payoff occurs, at which point the players will switch to the pure strategy by repeating the move they just made. Thus the number of players achieving a payoff rises geometrically at the rate of one half at each subsequent stage. If a meeting occurs, we might assume the players will coordinate in future by agreeing when to meet. Otherwise we suppose that players pick their coffee breaks as before. In that case, a meeting takes place with probability 1/10 on the first day, 9/100 the second day (9/10 times 1/10), 19/1000 the third day and so on. If there are N players who play an analogous game, an induction argument demonstrates that the probability of them spontaneously meeting together (in a one shot game) is 101-N. Now we change the structure of the game by giving one player, called the leader, power to send a message to the others proposing a meeting time. This immediately (and trivially) resolves the coordination problem, and establishes the value of coordination to the organization. It also illustrates the potential rent leaders can extract by reducing the coordination that takes place without their active involvement.

Focal points that distinguish themselves as more prominent solutions. Leadership: We define a leader as someone who chooses a pure strategy solution in a games where there are multiple pure strategy solutions. Note that leaders do not have an enforcement role, since by definition an equilibrium is self enforcing. In the examples we have reviewed on meetings, the coordination or leadership function is easy to play. We would not expect anyone to extract rents from performing this role because of competitive pressure to reduce the rent. However this need not be the case. Sometimes experience or skill is necessary to recognize potential gains to the players in the game.

Exercise *Accordingly, some laboratory exercises in this chapter focus on how well players coordinate their choices between different equilibrium.*

3.2 Conflicting Objectives

The examples of coordination games reviewed above illustrate the value of implicit or explicit agreements made before play begins. How such agreements are reached is the topic of the next chapter. In the previous examples it was easy to identify the set of coordinated strategic profiles. But recognizing possibilities for coordination are not always so evident. Let us consider the following example, this time as a finitely repeated game.

What happens when there are several equilibrium in the kernel game? We will see that the number of solutions in the repeated game increase dramatically. The kernel games we have studied above have multiple equilibrium.

Consider a partnership between two workers teamed together for N independent contracts and then disbanded. The partners individually and simultaneously choose between performing sloppy or diligent work. The N clients do not observe how hard each partner works but can assess the quality of the resulting product. Consequently each partner is paid the same wage, but that wage paid varies with the quality of the overall work. Since the quality of the work is evident, and each partner knows his own effort, he can also deduce how hard his partner worked. If both partners perform diligently, a client pays them $\$1,000$ each, while the effort and work opportunities forgone are valued at $\$400$. In this case each partner nets $\$600$. If one worker is diligent but his partner performs sloppy work, then this shows in the final product and the pair are only paid $\$800$ each. The diligent worker nets only $\$400$ in that case. However the worker who shirks reduces the opportunity cost of his time by completing other jobs and have more energy left for his pastime activities. The opportunity cost of performing a sloppy job is only $\$100$ rather than $\$400$. Consequently the worker who shirks nets $\$700$ providing his partner is diligent. If, however, both partners are sloppy, their poor output is more evident, and the client pays them only $\$400$ each, for a net reward of $\$300$. We model the partnership as a N fold repeated game for two players. The kernel is a simultaneous game with the same two actions for both players, work or shirk, illustrated in Figure 12.13.

Column Firm

Row firm	a	b
A	3, 3	1, 4
B	4, 1	0, 0

Figure 12.13
Nice and Nasty

First we solve the kernel game. There are no dominated strategies in this game, so the principle of iterated dominance does not apply. Note that working is a unique best response to shirking, and vice versa. Therefore there are two pure strategy Nash equilibrium, namely (work, shirk) and (shirk, work). There is also a mixed strategy equilibrium found by equating the payoffs from working with the payoffs from shirking, when the partner mixes. Let π denote the probability that a player works. Then his partner is indifferent between shirk and work if and only if

$$600\pi + 400(1-\pi) = 700\pi + 300(1-\pi) \Rightarrow \pi = 1/2$$

In multiple rounds of this game, the arguments about subgame perfection can be applied to show that any sequence of equilibria for the kernel game is part of an equilibrium strategy profile for the repeated game itself. An indication of how many equilibria there are can be given by showing which average payoff vectors are possible in equilibrium. that can arise from playing the kernel game a large but finite number of times. Simply repeating the pure strategy equilibrium produces (1,4) and (4,1), and the law of large numbers assures us that the average payoff from playing the mixed strategy equilibrium yields payoffs arbitrarily close to (8/3, 8/3). Supposing the players chose (large, small) for every even round of the game and (small, large) every odd round, their average payoff would be (5/2,5/2). Indeed any average payoff in the convex hull of the three kernel equilibrium points is asymptotically possible by judiciously choosing the proportion of times each of the equilibria are played. For example (pick a point in the interior and describe how to do it)

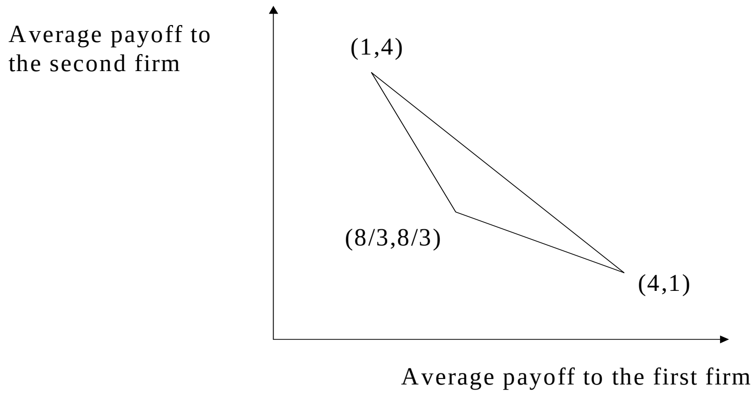

Figure 12.14

The limit set of average payoffs from repeatedly playing Nice and Nasty

The dots in Figure 12.14 show what average payoffs are possible in a three round game from playing different permutations of the solutions to the kernel game. As the number of rounds increase so do the number of dots. The convex hull in Figure 10.7 shows average payoffs that are possible from repeating equilibria of the kernel games indefinitely. There are, however, many other possible equilibrium outcomes that arise in the repeated game because the players choose dynamic strategies, that is making their choices contingent on play in previous rounds. For example suppose $N = 3,$ and the partnership played $(work, work)$ in the first round for $(600, 600)$, followed by $(work, shirk)$ yielding $(400, 700)$ and finally $(shirk, work)$ in the third round to obtain $(700, 400)$. This yields an average payoff of 567 to each player. Can this average payoff be attained by the team in equilibrium?

If this is a e second player deviates by forcing (1,4) in the first round, then both players settle for (4,1) in the two remaining rounds. Similarly if the first player deviates by forcing (4,1) in the first round, then both players agree to (1,4) in the remaining rounds. We now show that this is a self enforcing agreement, and therefore a pure strategy solution to the game.

Can Bond and Octopussy both earn more than 6 in a three period game? The outcome {(3,3), (1,4), (4,1)} comes from playing: {($bnice_1$, nice1), (nice2,nasty2), (nasty3, nice3)}. Is this history the outcome of a solution strategy profile to the 3 period repeated game?

Strategy for Bond

Round 1: nice1

Round 2: (…, nice1)⇒nice, otherwise nasty2

Round 3: (nasty1, …)⇒nice, otherwise nasty3

Bond should be nice in the first round. If Octopussy is nice in the first round, Bond should be nice in the second round too. If Octopussy is nasty in the first round, Bond should be nasty in the second. Bond should be nasty in the final round, unless he was nasty in the first round.

Strategy for Octopussy
Round 1: nice1
Round 2: (…, nasty1) ⇒ nice, otherwise nasty2
Round 3: (nasty1, …) ⇒ nasty3 otherwise nice3

Octopussy should be nice in the first round. Then if she followed her script in the first round, she should be nasty in the second. However if she forgot her lines in the first round and was nasty, then she should be nice in the second round. If Bond has was nasty in the first round, Octopussy should be nasty in the final round, but nice otherwise.

Verifying this strategy profile is a solution. Note that the last two periods of play, taken by themselves, are solutions to the kernel game, and therefore strategic form solutions for all sub-games starting in period 2. Checking for deviations by Bond in the first round To check whether being nice is a best response for James bond given that Octopussy chooses according to her prescribed strategy we compare:

1. (A_1, a_1) 3
2. Otherwise → (A_2, b_2) 1
3. Otherwise → (B_3, a_3) 4

8

With

1. (B_1, a_1) 4
2. Otherwise → (A_2, b_2) 1
3. (B_1, a_1) → (A_3, b_3) 1

6

Since 8 exceeds 6 row player does not profit from deviating in the first period. A similar result holds for the column player. Therefore the strategy is a SPNE. Establishing the $\{(3,3),(4.1),(1,4)\}$ is the outcome of a solution to this game, shows that average payoffs exceeding $2\frac{1}{2}$ can be achieved by both players in equilibrium.

Extending this example a little further, suppose the game lasts a finite number of periods, denoted by $N,$ and consider an extension to the agreement as follows: For the first $(N-2)$ periods, $(3,3)$ is played, the game ending with (4,1) and finally (1,4). If the second player deviates by forcing (1,4) in the first round, then both players settle for (4,1) in the two remaining rounds. Similarly if the first player deviates by forcing (4,1) in the first round, then both players agree to (1,4) in the remaining rounds. This argument shows that $(3-\frac{1}{N})$ can be achieved as an average payoff in the game. Following the arguments we made earlier, the interior of the triangle with vertices $(1,4)$, $(3,3)$ and $(4,1)$ can be reached in a solution to a finite horizon game

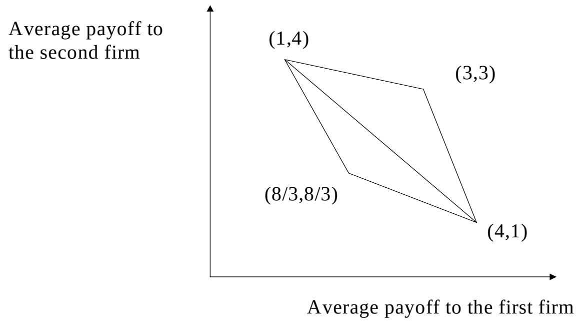

Figure 12.15

Average payoffs that are unattainable in a one period game

There is nothing optimal about self enforcing agreements. What is the lowest sum of payoffs in the 3 period repeated game? Unforgiven What is the lowest sum of payoffs in the 3 period repeated game that can be supported by a SPNE?

The outcome {(0,0), (1,4), (4,1)} is induced by playing {(nasty1, nasty1), (nice2,nasty2), (nasty3,nice3)} Can this outcome be supported by a SPNE? The outcome {(0,0), (1,4), (4,1)} is induced by playing {(B1, b1), (A2,b2), (B3,a3)} Strategy profiles supporting Unforgiven

Strategy for Clint Eastwood:

Round 1: nasty1

Round 2: (…, nice1) ⇒ nasty2 otherwise nice2

Round 3: (nice1, …) ⇒ nice3 otherwise nasty3

Strategy for the Gene Hackman:

Round 1: nasty1

Round 2: (…, nice1) ⇒ nice2 otherwise nasty2

Round 3: (nice1, …) ⇒ nasty3 otherwise nice3

Checking for a solution. Using the same methods as before one can show this is also a solution strategy profile for the three period game. More generally by punishing any deviation from the equilibrium path with the unfavorable kernel equilibrium repeated until the end of the game guarantees any payoff pair that averages more than the value given by individual rationality. Suppose the agreement for a three period game required both players choose $(0,0)$ in the first period, again followed by $(4,1)$ and then $(1,4)$, yielding an average payoff of $1\frac{2}{3}$. Using the same punishment strategy for any player who deviates, the same methods show that this is outcome of another solution to the game. Analogous to before this can be extended to show that $(1+\frac{1}{N})$ or something like that (to be fixed up) can be achieved. Figure 12.16 now fills in the area in the usual fashion.

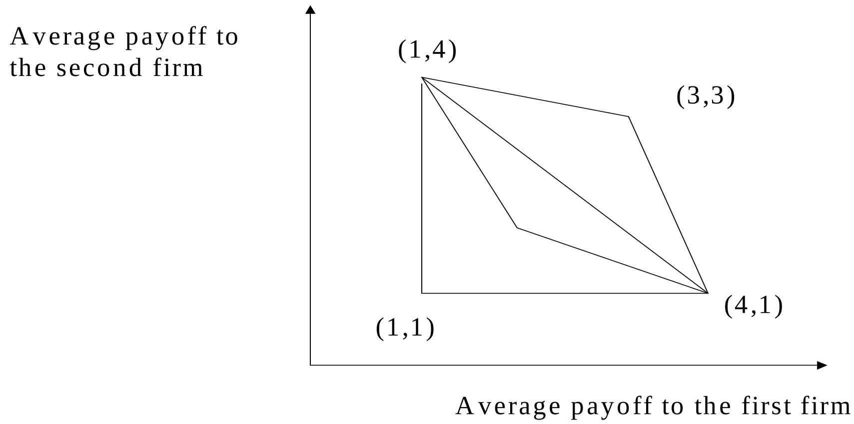

Figure 12.16

Equilibrium average payoffs in the long, finite horizon repeated game

Notice that no agreement can lead to a player receiving less than 1 because the player can unilaterally achieve that by playing on every round. further more nothing can be achieved outside of the $(1,4)$, $(3,3)$ and $(4,1)$ frontier, since there are no strategies that deliver that payoff. They are simply not technologically feasible in this game. Yet every other payoff can be achieved as the outcome to the repeated game if there are a sufficient number of rounds. A concise way of showing what can be achieved in two player finite horizon two player games is to first graph the payoffs that are feasible as in Figure 12.17.

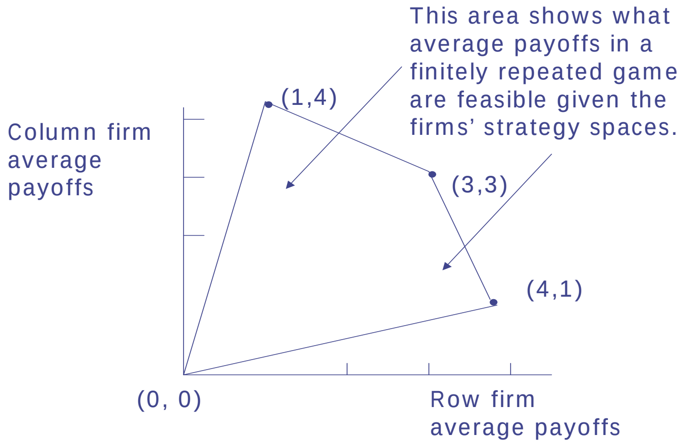

Figure 12.17

Feasible payoffs

Then we define individual rationality, and graph the individual rationality constraints. See Figure 12.18.

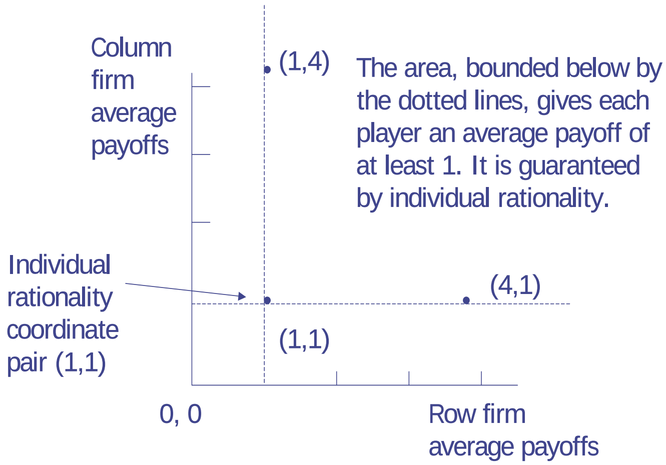

Figure 12.18
Individual rationality

Average payoffs in equilibrium are now found by taking the intersection of the two areas, as Figure 12.19 shows.

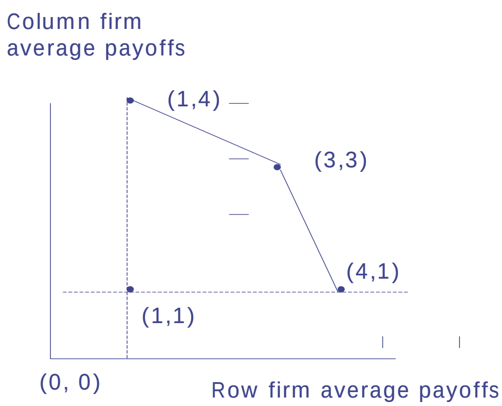

Figure 12.19
Average feasible payoffs meeting individuality constraints

The theorem in the next slide states that every pair in the enclosed area represents average payoffs obtained in a solution to the finitely repeated game. This theorem is a specialization of a result called the folk theorem, that extends to games with more than two players with minor modifications. It is an indeterminacy result, that (unfortunately) limits the predictive power of game theory when the kernel of the game supports multiple solutions. Folk theorem for two players:

Theorem *Let ω_1 be the worst payoff that Player 1 receives in a solution to the one period kernel game, let ω_2 be the worst payoff that Player 2 receives in a solution to the one period kernel game, and define $\omega = (\omega_1, \omega_2)$. In our example $\omega = (1, 1)$. Any point in the feasible set that has payoffs of at least ω can be attained as an average payoff to the solution of a repeated game with a finite number of rounds.*

4 Infinite Horizon Games with Discounting

Earlier in this chapter we claimed that the solution to a stage game is unique if the solution to each stage is unique and if there are only a finite number of rounds or periods. The folk theorem establishes that if there are multiple solutions to at least one of the stages, stage games support many more solutions than can be counted by the permuting the solutions of their component stages. The other factor limiting the applicability of the uniqueness result we derived is that the stage game have a finite number of stages. Infinite horizon games last an indefinite number of periods, either because they simply never end, or because at the end of each period, there is a strictly positive probability of continuing the next period. When players realize that their relationship does not have a foreseeable terminal node, new possibilities for cooperation and mutual benefit emerge. This section relaxes the other qualification of the uniqueness result, by analyzing games where each stage has a unique solution, but there are an infinite number of rounds. We now show how and when cooperative behavior between group members can be enforced despite their individually conflicting objectives, by credibly threatening strategies that punish actions which harm a collective interest.

4.1 Perfect Monitoring

To show the main result in this section we consider an infinite horizon repeated game with discounting between two players who move simultaneously each period and observe the outcome at the end of the period. We label the players as General Motors and Ford, and imagine that at one point in history these two car manufacturers accounted for most of the automobile sales in the North American market. For simplicity we assume that dealers are responsible to the directions from central management, and there are only two prices at which either manufacturer would consider selling its cars, cheap and expensive. The net profits (per period) from marketing cars at the lower price are greater than choosing the higher price, regardless of what marketing strategy the other firm chooses. Consequently selling cars cheaply is a dominant strategy. Hence there is a unique solution in the one stage game for each firm to charge the low price. Furthermore the uniqueness theorem we derived earlier in this chapter implies that both Ford and G.M. would price their cars inexpensively every period in a repeated game with a finite horizon.

In addition to the assumptions made above that uniquely determine the solution to a finitely repeated game, suppose the profits to both firms from charging high prices are greater than if both firms charge low prices. This additional assumption implies that within each stage the configuration of payoff inequalities mimics the prisoner's dilemma.

To derive the solutions to this infinite horizon game with discounting, we now introduce some notation for analyzing the game. Denote by $\{e_{nt}, c_{nt}\}$ the choice set for the n^{th} player in period $t,$ where e_{1t} stands for Ford selling expensive cars in the t^{th} period, c_{2t} means General Motors sells cheap cars in that period and so forth. Let $s_t \equiv (s_{1t}, s_{2t}) \in \{e_{nt}, c_{nt}\} \times \{e_{nt}, c_{nt}\}$ denote the respective choices made by Ford and

General Motors in period t. A strategy for Ford is a complete set of plans about what to charge in each period, a function of all the choices that have been made by both players in the past. Thus s_{1t} can potentially depend on the history of pricing choices $\{s_{1\tau}, s_{2\tau}\}_{\tau=1}^{t-1}$ made up until and including period $t-1$.

We denote by $\pi_n(s_t)$ the current profits of the n^{th} player in period t as a function of the players' choices made at that time, and assume that current profits are converted to present value units with a constant interest rate, denoted by r. Defining the discount factor as $\beta = (1+r)^{-1}$, the present value of profits Ford obtains in period t may then be expressed as $\beta^t \pi_1(s_t)$, and a similar expression, $\beta^t \pi_2(s_t)$, holds for General Motors. The objective function for the n^{th} player is

$$\sum_{t=0}^{T} \beta^t \pi_n(s_t)$$

where $T = \infty$ in this application. Note that in this formulation of the game we have not allowed for the possibility that future payoffs might be uncertain, because nature has no direct role in the game, and we shall restrict ourselves to considering pure strategy equilibrium solutions.

The assumption that marketing cheap cars is a dominant strategy for Ford in a single period game played at time t is captured by the inequalities

$$\pi_1(c_{1t}, s_{2t}) > \pi_1(e_{1t}, s_{2t})$$

for $s_{2t} \in \{e_{2t}, c_{2t}\}$. Similarly marketing cheap cars is a dominant strategy for General Motors in a single period game played at time t if and only if

$$\pi_2(s_{1t}, c_{2t}) > \pi_2(s_{1t}, e_{2t})$$

for $s_{1t} \in \{e_{2t}, c_{2t}\}$. Finally both firms make more profits when they both sell expensive cars compared to when they both sell cheap cars corresponds to the pair of inequalities

$$\pi_n(e_{1t}, e_{2t}) > \pi_n(c_{1t}, c_{2t})$$

for $n \in \{1, 2\}$. Table 12.20 depicts the one period simultaneous move game upon which the repeated game is based.

		Column Player	
		ht	lt
Row Player	Ht	R, R	d, D
	Lt	D, d	r, r

Figure 12.20

The strategic form of a kernel game

In the finite horizon case $T < \infty$, and the arguments we discussed for finite horizon games extends in a simple way to this class of games. Since (c_{1t}, c_{2t}) is the unique solution to a single stage game played at t, the solution to a finite horizon game formed from repeating the stage is to play it. Thus the unique solution to this game is:

$$\{c_{1t}, c_{2t}\}_{t=0}^{T}$$

It is straightforward to check that this strategy profile is also a solution to the infinite horizon case (when $T = \infty$). Are there any other solutions?

We reconsider a class of strategies that we introduced in our analysis of multistage finite horizon games, called trigger strategies. These strategies support solutions where strategic investment occurs. Suppose Ford markets expensive cars unless or until General Motors offers a cheap car, at which point it offers only sells cheap cars thereafter. The sequence of periods during which only expensive cars are marketed is called the cooperative phase, while the sequence of periods during which both firms sell only cheap cars is called the punishment phase. In terms of the notation we have already developed suppose:

$$s_{1t} = \begin{cases} e_{10} \text{ if } t = 0 \\ e_{1t} \text{ if } s_{\tau} = (e_{1\tau}, e_{2\tau}) \text{ for all } 0 \le \tau < t \\ c_{1t} \text{ otherwise} \end{cases}$$

Regarding General Motors, let us similarly assume

$$s_{2t} = \begin{cases} e_{20} \text{ if } t = 0 \\ e_{2t} \text{ if } s_{\tau} = (e_{1\tau}, e_{2\tau}) \text{ for all } 0 \le \tau < t \\ c_{2t} \text{ otherwise} \end{cases}$$

Are trigger strategies a solution to this game? To determine whether the trigger strategies are a solution, we only need to check whether the subgames are solved by them. There are two kinds of subgames, depending on whether somebody has cheated in the past or not. Initially let us assume that Ford and/or General Motors has offered a cheap car at some past period $\tau < t$. In that case the trigger strategy profile is in the punishment phase and requires both manufacturers to market cheap cars. In this case the value to Ford of both players following the trigger strategy from period t onwards is thus

$$\sum\nolimits_{\rho=t}^{\infty} \beta^{\rho} \pi_1(c_{1\rho}, c_{2\rho}) = \beta^{t} \pi_1(c_{1t}, c_{2t}) + \sum\nolimits_{\rho=t+1}^{\infty} \beta^{\rho} \pi_1(c_{1\rho}, c_{2\rho})$$

whereas if Ford deviated from the trigger strategy by marketing an expensive car in period its value is

$$\beta^{t} \pi_1(e_{1t}, c_{2t}) + \sum\nolimits_{\rho=t+1}^{\infty} \beta^{\rho} \pi_1(c_{1\rho}, c_{2\rho})$$

The loss to Ford from deviating by marketing an expensive car in period t is

$$\beta^t[\pi_1(e_{1t},c_{2t}) - \pi_1(c_{1t},c_{2t})] < 0$$

A similar calculation applies to General Motors. It immediately follows that the punishment phase of the trigger strategy solves subgames in which a cheap car has been marketed in the past. Another of putting this is to say that the punishment phase is self enforcing.

Now consider game histories leading up to period t in which no cheap car has ever been marketed by either firm, including the first period. If both firms follow the trigger strategy, they will market expensive cars in period $t,$ so next period the history will be $\{(e_{1\tau},e_{2\tau})\}_{\tau=0}^{t}$. If they continue following the trigger strategy in periods $\rho > t$, the history will be successively updated to $\{(e_{1\tau},e_{2\tau})\}_{\tau=0}^{\rho-1}$ and the firms will never market a cheap car. In other words the cooperative strategy is self sustaining. Therefore at period t in a cooperative phase the value to General Motors from both manufacturers following the trigger strategy is

$$\sum_{\rho=t}^{\infty}\beta^{\rho}\pi_2(e_{1\rho},e_{2\rho}) = \beta^t\pi_2(e_{1t},e_{2t}) + \sum_{\rho=t+1}^{\infty}\beta^{\rho}\pi_2(e_{1\rho},e_{2\rho})$$

If, however, General Motors defects from this strategy during the cooperative phase in period t by marketing a cheap car, and both firms follow the trigger strategy in all other respects, then Ford would market an expensive car in period $t,$ but deviating will be instigate the punishment phase in period $t+1,$ never to end. Thus both firms would market cheap cars from period $t+1$ onwards, and at period $\rho > t$ the history of car marketing will be

$$\{(e_{10},e_{20}),\ldots,(e_{1,t-1},e_{2,t-1}),(e_{1t},c_{2t}),(c_{1,t+1},c_{2,t+1}),\ldots,(e_{1,\rho-1},e_{2,\rho-1})\}$$

Therefore the value to General Motors of defecting from the trigger strategy if is in the cooperative phase is

$$\beta^t\pi_2(e_{1t},c_{2t}) + \sum_{\rho=t+1}^{\infty}\beta^{\rho}\pi_2(c_{1\rho},c_{2\rho})$$

Taking the difference of these expressions, the net value for General Motors of deviating from the trigger strategy when it is in the cooperative phase at period t is

$$\beta^t\pi_2(e_{1t},c_{2t}) + \sum_{\rho=t+1}^{\infty}\beta^{\rho}\pi_2(c_{1\rho},c_{2\rho}) - \beta^t\pi_2(e_{1t},e_{2t}) - \sum_{\rho=t+1}^{\infty}\beta^{\rho}\pi_2(e_{1\rho},e_{2\rho})$$

Factoring out and combining the terms in the summation we obtain

$$\beta^t\left\{\pi_2(e_{1t},c_{2t}) - \pi_2(e_{1t},e_{2t}) + \sum_{\rho=1}^{\infty}\beta^{\rho}[\pi_2(c_{1\rho},c_{2\rho}) - \pi_2(e_{1\rho},e_{2\rho})]\right\}$$

Now applying the formula for summing a geometric series yields

$$\beta^t\left\{\pi_2(e_{1t},c_{2t}) - \pi_2(e_{1t},e_{2t}) + \frac{\beta}{1-\beta}\pi_2(c_{1\rho},c_{2\rho}) - \pi_2(e_{1\rho},e_{2\rho})\right\}$$

This expression has an intuitive interpretation. The current benefit to General Motors from defecting in the cooperative phase is

$$\pi_2(e_{1t},c_{2t}) - \pi_2(e_{1t},e_{2t})$$

while the cost is the present value from remaining in the punishment phase compared to the profits that would have been made if the firms had been made in the cooperative phase, which is the product of loss per period, $\pi_2(c_{1\rho}, c_{2\rho}) - \pi_2(e_{1\rho}, e_{2\rho})$, and the value of an annuity of one unit $(1-\beta)^{-1}$. In order for General Motors to break the trigger strategy, the immediate gain to them from defecting and selling cheap cars while Ford is still selling expensive ones must exceed the punishment of lower profits that inevitably follows. Again a analogous expression can be derived for Ford. A necessary and sufficient condition for a trigger strategy to solve this game is that neither firm benefits from unilaterally deviating in the cooperative phase.

Before embarking on a numerical analysis, we note that the key to the inequalities are the ratios of the payoffs. In an empirical application it may be useful to normalize by one of the payoffs, say $\pi_n(e_{1t}, e_{2t})$, and proceed in those terms.

4.2 A numerical example

To help develop a quantitative sense of the factors involved, suppose the two companies are symmetric to each, meaning that they earn the same payoffs when they are in the same relative position. That is

$$\pi_1(e_{1t}, e_{2t}) = \pi_2(e_{1t}, e_{2t})$$
$$\pi_1(c_{1t}, c_{2t}) = \pi_2(c_{1t}, c_{2t})$$
$$\pi_1(e_{1t}, c_{2t}) = \pi_2(c_{1t}, e_{2t})$$
$$\pi_1(c_{1t}, e_{2t}) = \pi_2(e_{1t}, c_{2t})$$

Table 10.8 illustrates the strategic form of the kernel game when payoffs are symmetric. In this example

$$\pi_1(e_{1t}, e_{2t}) = \pi_2(e_{1t}, e_{2t}) = 10$$
$$\pi_1(c_{1t}, c_{2t}) = \pi_2(c_{1t}, c_{2t}) = 5$$
$$\pi_1(e_{1t}, c_{2t}) = \pi_2(c_{1t}, e_{2t}) = 20$$
$$\pi_1(c_{1t}, e_{2t}) = \pi_2(e_{1t}, c_{2t}) = 0$$

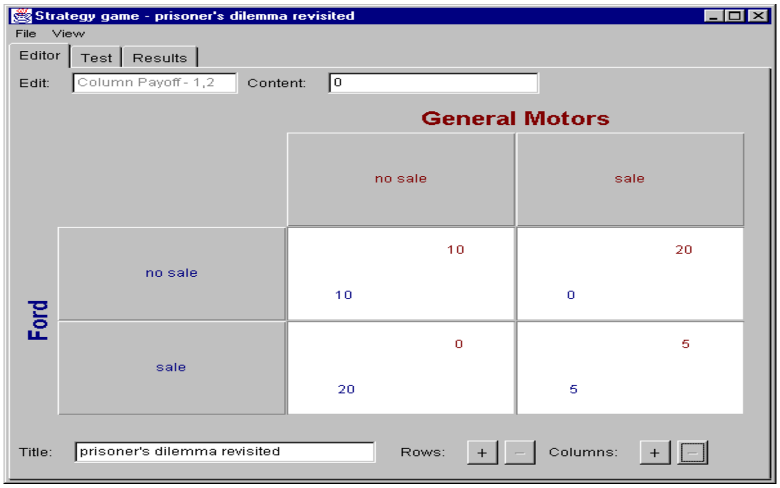

Figure 12.21

Suppose the gestation period between the time a player defects and the period the other player can retool and produce an inexpensive car is between two and three years. This might seem too short but for the fact that industry intelligence is reveal to a rival what a firm's intentions are before the product is introduced to the marketplace. According we suppose initial that the interest rate is 25 percent, implying $\beta = 0.8$. Then the value to each firm from sustaining the cooperative phase is

$$\begin{aligned} &\pi_1(e_{1t}, e_{2t})\left(\frac{1+r}{r}\right) \\ &= 10\left(\frac{1.25}{0.25}\right) \\ &= 50 \end{aligned}$$

and the value of each firm upon reaching the punishment phase is

$$\begin{aligned} &\pi_1(c_{1t}, c_{2t})\left(\frac{1+r}{r}\right) \\ &= 5\left(\frac{1.25}{0.25}\right) \\ &= 25 \end{aligned}$$

Therefore the value to each firm of defecting from the cooperative phase is

$$\begin{aligned} &\pi_1(c_{1t}, e_{2t}) + \pi_1(c_{1t}, c_{2t})\left(\frac{1}{1+r}\right)\left(\frac{1+r}{r}\right) \\ &= \pi_1(c_{1t}, e_{2t}) + \pi_1(c_{1t}, c_{2t})\left(\frac{1}{r}\right) \\ &= 20 + 5\left(\frac{1}{0.25}\right) \\ &= 40 \end{aligned}$$

It follows that cooperation with both firms marketing an expensive car is a solution to this repeated game.

4.3 Factors determining cooperation

The arguments we have used suggest that in frameworks of the sort described above, cooperation can be introduced at the very beginning of the game, and be sustained by a credible threat that is never imposed throughout the course of play, or that cooperation cannot be achieved at all. We remark that even if the parameters of the problem permit cooperation, there is still typically scope for intermittent cooperation and reversion to punishment strategies. As we found in finite round games with multiple solutions, there are many other solutions to the infinite horizon game when the cooperative solution is attainable. However coordinated players have an incentive to avoid such solutions, if they can reach an agreement about the gains from trade. Consequently the framework does not provide a satisfactory explanation of why cooperation is sometimes an intermittent phenomenon.

Our discussion highlights three factors that determine whether a relationship between strategic partners is cooperative or adversaries. To recapitulate, the three factors are:

1. the gains of maintaining, compared to the benefit of defecting from, a cooperative arrangement

2. the losses from destroying a cooperative arrangement and reverting to an adversaries relationship

3. the duration of the relationship measured in discounted time (to reflect the probability of survival and the interest rate).

In the applications above, one of two possibilities emerges. Either cooperation can be maintained for the whole game, and the credible punishment threat is never administered, or cooperation can never be achieved. We can extend this result. Suppose the three factors described above are known for all future periods. Then either cooperation is established at some point and maintained thereafter or cooperation is never established. To summarize, cooperation never breaks down when all the factors are known in advance and the players are rational. This is what characterizes stable relationships.

There are several reasons why cooperative relationships can break down or revert to adversaries confrontations:

1. When the activities of one or more players cannot be monitored

2. When one or more of the three factors described above is a time dependent stochastic process

3. If some of the payoff relevant information to one player is hidden from the other one.

Such situations may be described as volatile relationships, and we now seek to analyze them.

4.4 Time dependent factors

A second reason for volatility in relationships is that the costs and benefits of cooperation may vary stochastically over time. Note that if there was a fixed point in time when both parties knew that the conditions for cooperation would not be met, then the arguments we applied in the finite horizon game would apply, and the unique solution to the game would preclude any cooperation.

We now suppose there is some chance of fluctuating between regimes where the benefits of cooperating are high to a regime were those benefits are minimal. In the example displayed below, automobile demand is sometimes high and sometimes low. After each round is played the stage game ends with probability of 0.1, and probability of 0.9 of continuing at least one more round. If automobile demand is high in the current period, the game continues with probability 0.7 into another round of high demand, and with probability 0.2 into a round where demand is low. Once in a low demand state of demand, it a little more likely to stay there than return to the high state, the odds being 5 to 4. In both stages of this game, offering a rebate is a strictly dominant strategy for both companies. However the immediate gain of defecting from a cooperative strategy where neither company ever offers rebates is greater in the low demand state than the higher one, 7 versus 1.

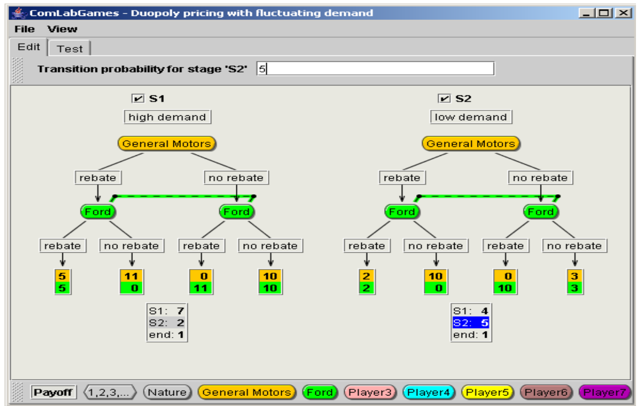

Figure 12.22

How much cooperation can be sustained? It is easy to check that if high demand prevailed every period, that is the stage transition weights were $(9, 0, 1)$ instead of $(7, 2, 1)$, then a grim trigger strategy would support the cooperative solution outcome of no rebates, whereas if low demand prevailed every period, and the transition weights were $(0, 9, 1)$ instead of $(4, 5, 1)$, then the unique solution would induce both firms to offer rebates every period.

To check whether cooperation can be sustained throughout both stages we first calculate the value of each firm when neither deviates, and then compute the value to a firm from deviating by offering a rebate at either stage when the maximal penalty is imposed. If the value of the latter quantity exceeds the value of the former quantity, we conclude that cooperation throughout the whole game is unsustainable. We denote by (v_{11}, v_{21}) the value of each firm in the two states when neither ever offers rebates. These values must satisfy the recursion:

$$\begin{pmatrix} v_{11} \\ v_{21} \end{pmatrix} = \begin{pmatrix} 10 \\ 3 \end{pmatrix} + 0.8 \begin{bmatrix} 0.75 & 0.25 \\ 0.5 & 0.5 \end{bmatrix} \begin{pmatrix} v_{11} \\ v_{21} \end{pmatrix}$$

$$= \begin{bmatrix} 1 - (0.8)(0.75) & -(0.8)(0.25) \\ -(0.8)(0.5) & 1 - (0.8)(0.5) \end{bmatrix}^{-1} \begin{pmatrix} 10 \\ 3 \end{pmatrix}$$

$$= 10 \begin{bmatrix} 4 & -2 \\ -4 & 6 \end{bmatrix}^{-1} \begin{pmatrix} 10 \\ 3 \end{pmatrix}$$

$$= \frac{10}{28} \begin{bmatrix} 6 & 2 \\ 4 & 4 \end{bmatrix} \begin{pmatrix} 10 \\ 3 \end{pmatrix}$$

The top line of the first equality can be explained as follows. The value of being in the high demand state when neither firm ever offers rebates, v_{11}, is current profits, 10, plus the remaining value of the firm next period discounted by the probability of continuation, 0.8. The value of the firm next period is either v_{11}, which occurs with probability 0.75, or v_{21}, which occurs with probability 0.25. The bottom of line of the first equation can be justified in a similar manner. Finally the second equality is found by making the vector (v_{11}, v_{21}) the subject of the equation. Upon inverting the matrix in the third equality we obtain the solution $(v_{11}, v_{21}) =$. The same technique can be applied to calculate the value of the game from playing the dominant strategy for the kernel games. We solve for (v_{12}, v_{22}), the value of each firm when both always offer rebates, by substituting $(10, 3)$ with $(5, 2)$ and (v_{12}, v_{22}) for (v_{11}, v_{21}) in the formula above to obtain $(v_{12}, v_{22}) =$.We are now in a position to evaluate whether collusive behavior can be sustained throughout both phases. The value from adhering to the collusive agreement is $(v_{11}, v_{21}) = (\quad)$ (depending on which state of demand holds) but the value from deviating is $(20 + 0.8v_{12}, 10 + 0.8v_{22}) = (\quad)$.

Since the gains from adhering to a tacit agreement of no rebates are relatively small when there is low demand, it might is reasonable to check whether collusive behavior could be enforced when there is a high state of demand, but not the low one. If this degree of collusion could be sustained then current profits would be 10 in periods of high demand, but only 2 in periods of low demand. We let (v_{13}, v_{23}) denote the values of the firm in the two states under this arrangement. can be directly applied here too. To solve for (v_{13}, v_{23}) we can directly apply the algebraic manipulation described in the previous paragraph. Substituting (v_{13}, v_{23}) for (v_{11}, v_{21}) and $(10, 2)$ for $(10, 3)$ in the recursion for firm value yields

$$\begin{pmatrix} v_{13} \\ v_{23} \end{pmatrix} = \begin{bmatrix} 0.75 & 0.25 \\ 0.5 & 0.5 \end{bmatrix} \begin{pmatrix} 10 \\ 2 \end{pmatrix}$$

In this case deviant behavior would not be observed in the low demand periods since it entails lower current profits and a lower expected value of the firm next period.

5 Incomplete Information in Markov Games

The last issue we address on this topic is the role of incomplete information.

5.1 Imperfect monitoring

The situation becomes more complicated when players in a repeated or stage game cannot deduce the choices made at a given stage from its outcome. One rationale for why the price of crude oil oscillates is that OPEC members cannot infer exactly how whether each country is abiding to its allotment as determined by the OPEC cartel. If only aggregate oil production by OPEC members and nonmembers is common knowledge, It is conceivable that OPEC might make its allocation decisions and recommendations to its members on the basis of that aggregate.

The following experiment captures the elements of the dilemma facing a cartel as its seeks to enforce collusive behavior. There are two (or more) cartel members who choose either high or low production each period, and we denote their choices in period $t \in \{0,1,\ldots\}$ by (q_{1t}, q_{2t}) where $q_{jt} \in \{l,h\}$. The current profits to each producer depends on the choices of both, and also crude oil supply by producers who are not cartel members, and are not modeled as players in this experiment. At the end of period t each cartel producer receives his own payoff, as well as information about total output produced that period, Q_t, which takes on two values, denoted H and L respectively. The profits of each producer $j \in \{1,2\}$ only depends on his own production choice and the signal received about total production. We denote the current profit of the producer $j \in \{1,2\}$ in period t by $\pi(q_{jt}, Q_t)$ and assume:

$$\pi(h,L) > \pi(l,L) > \pi(h,H) > \pi(l,H)$$

We assume that both producers have a sufficiently large share of the market that their individual choices significantly affect the total output. Let $\Pr(q_{1t}, q_{2t})$ denote the probability that $Q_t = H$ given choices (q_{1t}, q_{2t}). Naturally, the higher the output of either producer, the higher the probability that high aggregate output is recorded. That is

$$\Pr(h,h) > \Pr(h,l) = \Pr(l,h) > \Pr(l,l) > 0$$

Thus the expected current profit from choosing h when the other producer chooses l is

$$E[\pi(h,Q)|q_{2t} = l] \equiv [1 - \Pr(h,l)]\pi(h,L) + \Pr(h,l)\pi(h,H)$$

We assume that choosing h yields higher expected profits at the end of the period regardless what the other player decides. This assumption can be expressed as

$$\begin{aligned} E[\pi(h,Q)|q_{2t}] &\equiv [1 - \Pr(h,q_{2t})]\pi(h,L) + \Pr(h,q_{2t})\pi(h,H) \\ &> E[\pi(l,Q)|q_{2t}] \equiv [1 - \Pr(l,q_{2t})]\pi(l,L) + \Pr(l,q_{2t})\pi(l,H) \end{aligned}$$

for $q_{2t} \in \{l,h\}$. This assumption implies that absent repercussions, the dominant strategy of each cartel member is to choose h each period. The payoff to producers at the end of the game are the discounted sum of payoffs

$$\sum_{t=0}^{\infty} \beta^t \pi(q_{jt}, Q_t)$$

where $\beta \in (0,1)$ is a discount factor that reflects the interest rate and the probability of the game continuing another period.

The framework and its accompanying assumptions imply there is a subgame perfect equilibrium solution in which each cartel producer picks the strategy $q_{jt} = h$ for all $t \in \{0,1,\ldots\}$ and $j \in \{1,2\}$. It is also easy to infer from our discussion of the perfect monitoring case that the trigger strategy profile there is also a solution to the game if the probability of reaching the high aggregate demand when both players choose low production (denoted by $\Pr(l,l)$), the gains from deviating (measured by strength of the inequality above), and the interest rate (which is $\beta^{-1} - 1$) are all small enough. From the cartel's perspective, however, this equilibrium has the unfortunate feature, that the reversionary state is absorbing: once in the punishment phase, they cannot escape.

Hence the question naturally arises whether there exists a less severe punishment that nevertheless deters either party from defecting from the cooperative choice of low production.

We seek symmetric solutions where both firms produce low output during cooperative phases of the game, where a finite sequence of high output choices by both firms characterize punishment phases, and where punishment phases are triggered by high aggregate outputs occurring too frequently. Let V_1 denote the current value of the firm next period if low aggregate output is observed and V_2 the current value next period if a reversionary phase has just begun. By definition

$$V_2 = E\left[\sum_{t=0}^{\tau-1} \beta^t \pi(q_{jt}, Q_t) | (q_{1t}, q_{2t}) = (l, l)\right] + \beta^\tau V_1$$

where τ is the length of the reversionary phase.

For simplicity we focus on the existence of solutions where punishment is triggered as soon as high aggregate output is observed. To deter either firm from deviating in a cooperative phase we require

$$[1 - \Pr(l,l)][\pi(l,L) + \beta V_1] + \Pr(l,l)[\pi(l,H) + \beta V_2] > [1 - \Pr(h,l)][\pi(h,L) + \beta V_1] + \Pr(h,l)[\pi(h,H) + \beta V_2]$$

The left side of the inequality is the expected value of the first producer if both producers choose $l.$ Given $(q_{1t}, q_{2t}) = (l,l),$ the probability of observing aggregate production of H and consequently reverting to the punishment phase is $\Pr(l,l),$ reducing the expected value of both firms to $[\pi(l,H) + \beta V_2]$ each. Otherwise the value of both firms is $[\pi(l,L) + \beta V_1],$ reached with probability $[1 - \Pr(l,l)].$ On the right side is the expected value of the first producer when he chooses $h,$ and the other producer chooses $l.$ Conditional on aggregate production $Q_t \in \{H, L\},$ the first producer is better off choosing h over $l,$ because $\pi(l,L) > \pi(h,L)$ and $\pi(l,H) > \pi(h,H).$ Since $\Pr(l,l) < \Pr(h,l),$the inequality can only hold if the difference between V_1 and V_2 is sufficiently large: that is the feasibility of a trigger strategy solution rests on the fact that deviating from l to h is more likely induce a punishment phase, which might deter a producer if it is costly enough.

We now further simplify this problem by assuming that

$$0 < \Pr(l,l) < \Pr(h,l) = \Pr(l,h) = 1$$

This assumption implies that whenever either firm deviates from the cooperative strategy of low production, the aggregate indicator of production is $H,$ but that the indicator is not a perfect monitoring device, since $\Pr(l,l) > 0,$ meaning that even when both firms fully adhere to the cooperative strategy, the indicator of aggregate production is sometimes $H.$

In this specialization the inequality for incentive compatibility to cooperation simplifies to

$$[1 - \Pr(l,l)][\pi(l,L) + \beta V_1] + \Pr(l,l)[\pi(l,H) + \beta V_2] > \pi(h,H) + \beta V_2$$

and the value of the firm at the beginning of a punishment phase becomes:

$$V_2 = \sum_{t=0}^{\tau-1} \beta^t + \beta^\tau V_1 = \pi(l,H)(1-\beta^\tau)/(1-\beta) + \beta^\tau V_1$$

words, the expected current benefit from deviating by setting a high production level is less than the product of a successful deterrent and the difference in the probabilities induced by producing at a high versus low level.

The most severe deterrent is to impose an eternal punishment, which would imply

$$V_2 = \pi(l,H)/(1-\beta)$$

substituting this equation back into the incentive compatibility condition we obtain

$$[1-\Pr(l,l)][\pi(l,L)+\beta V_1] + \Pr(l,l)\pi(l,H) > \pi(h,H) + [1-\Pr(l,l)]\beta\pi(l,H)/(1-\beta)$$

Making V_1 the subject of the inequality we obtain

$$[1-\Pr(l,l)][\pi(l,L)+\beta V_1] - [1-\Pr(l,l)]\pi(l,H) > \pi(h,H) - \pi(l,H) + [1-\Pr(l,l)]\beta\pi(l,H)/(1-\beta)$$

$$\beta V_1 > \frac{\pi(h,H)-\pi(l,H)}{[1-\Pr(l,l)]} - \pi(l,L) + \frac{\pi(l,H)}{(1-\beta)}$$

We remark that this inequality simplifies to the model of perfect monitoring when $\Pr(l,l) = 0$.

There are two further extensions we should mention. The first is whether the overall value of the collusive arrangement can be increased by reducing the length of the punishment phase. This is a straightforward exercise which we leave to the reader. Second is the question about whether the firms should revert to a punishment phase as soon as high aggregate output is observed, or whether some lenient criterion should be impose instead as an acknowledgement that high output is observed even when both firms fully adhere to the cooperative arrangement. This is a difficult problem to solve, and lies beyond the scope of this course.

5.2 Hidden information

The third reason for volatility in relationships that we investigate revolves around private information that players have about themselves. Consider a workforce pool of laborers who team up in pairs to undertake construction jobs throughout the suburban neighborhoods in a big city. Workers are initially assigned at random to each other for their first job. Jobs are identical. Once assigned to his job, each partner chooses simultaneously whether to fully complete the job or only perform shoddy work. The payment each partner receives at the end of the period depends on the joint performance of the pair. At the end of every period, they decide whether to continue working with each other or not. The partnership is dissolved unless both partners wish to remain paired to each other. Teams who decide to break their match after working with each other join the pool of unmatched workers, and are then randomly assigned new partners. Workers differ in their ability, skilled workers are able to perform competent work with less effort than unskilled workers. However it is impossible to distinguish between the two types, because unskilled workers can also perform competently, albeit at greater effort.

If two skilled workers are teamed together then their payoff matrix looks like:

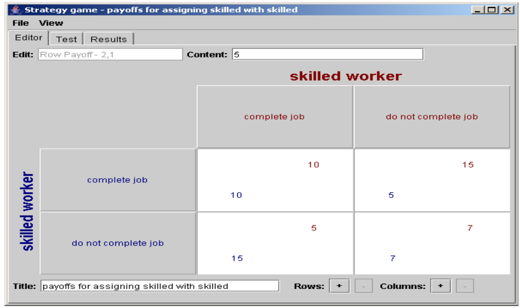

Figure 12.23

If a skilled worker is matched with an unskilled worker, the matrix looks like:

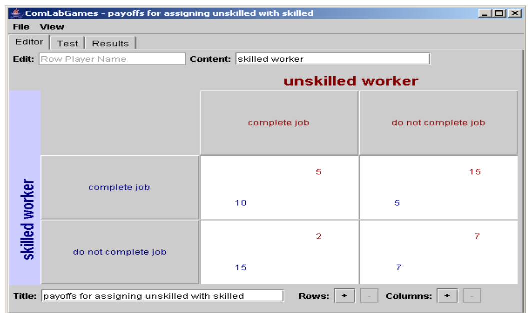

Figure 12.24

Finally if two unskilled workers are matched together then the payoff matrix looks like:

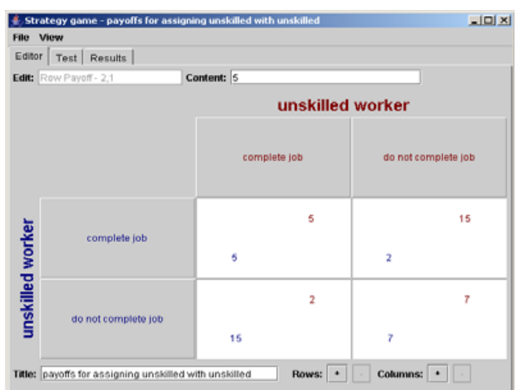

Figure 12.25

We assume the discount factor as the probability of continuing work next period,and in this way consider a game comprising an indefinite number of rounds. Regardless of who meets whom, the dominant strategy for a one round game is to leave the job incomplete.

Consider first the full information case where players are initially randomly assigned to match pairs, play a single round, and then observe their partner's type before making a decision about dissolving the match. It is optimal for a pair of skilled workers to reach a cooperative agreement to complete successive jobs once they recognize their partner is skilled, providing the probability of continuing the partnership is high enough to support a trigger strategy. This cooperative strategy cannot be supported by any other pair types. Therefore one solution to the game is for everyone to quit their first job before completion, and upon discovering their partner's type vote to continue the partnership if he is skilled. Thus skilled partners remain together and enter a cooperative partnership of completing jobs, while everyone else returns to the general pool of workers which is gradually depleted of skilled workers.

If there are a sufficiently low proportion of unskilled workers in the total population, there is also a solution in which skilled workers complete the job in the first period, and retain their partner if and only if the partner does the same thing. This solution does not use the worker's type as information determining the continuation of the match. So if this strategy profile is a solution for the full information case, it also holds in the hidden information case too. Completing the job in the first period sends a signal to your partner that you are a skilled worker.

6 Summary

This chapter introduced the dual concepts of stages and transition probabilities as tools for parsimoniously representing games where there are elements of repetition. One question we answered is whether solutions to stage and Markov games can be found by simply joining the solutions for each stage played independently. In stage games (where the probability transition is independent of the choices made), strategy profiles formed this way are indeed a subset of the solution set for the whole game. In Markov games (where the probability transitions are affected by choices), the arguments developed in Chapter 4 on investment and Chapter 6 on perfect information games apply here with equal force: we would not expect this approach generate a solution to the game.

A second question we answered is whether all the solution to stage games can be derived by piecing together the solutions of each stage game. to one of the solutions to the game can indeed be derived this way. Here we focused on repeated games, games with one stage that are repeated, either a finite number of times, or indefinitely with some probability of continuing from one round to the next. The first result is that if every stage has a unique solution, and the stage game is finite, then it also has unique solution, found by piecing the stage solutions together. However if more than one of the stages has multiple solutions, or a stage with multiple stages, then there are solutions to the game which cannot be found by piecing together the stages.

The third result is called the folk theorem, applies to finitely repeated games in which the kernel game supports multiple strategic solutions. We characterized the multiplicity of solutions to the repeated game in this case. If a kernel game is uniquely solved, there is a unique solution to a game that repeats the kernel a finite number of times. However there may be multiple solutions if the kernel is repeated indefinitely. The fourth result, which applies to infinite horizon games with discounting, establishes the existence of strategic form solutions that lack a counterpart in analogous games with the same kernel but a fixed finite horizon. We focused on one type of strategy, called a trigger strategy because it rationalizes the notion of collective punishment for those who deviate from agreed upon norms. In a trigger strategy solution, playing what would be the best reply in the kernel yields less than the long term benefits obtained by cooperating with the other players. Opportunities for coordination depend on the payoff parameters and the probability of repetition (or the discount factor.)

Thus if the stage in a two player repeated game has multiple solutions, then the area enclosed by the payoffs and the individual rationality constraints determines the set of average payoffs that can be attained. The solution strategies may induce individual players to engage in investment or commitment even though there is nothing in the technology to

Taken at face value both these results are troublesome for game theory because they reduce the power of the theory to deliver testable predictions. Furthermore, the coordination issues involved in arriving at one of several strategic form solutions seem incredible, revealing another unattractive feature of the Nash equilibrium concept. The fact that so many alternatives are available naturally raises the question about how players agree upon any one of the solutions. This led us to a new interpretation of multiple pure strategy solutions, as self enforcing agreements. Leaders choose amongst multiple solutions to achieve coordination between players. The less the potential for coordination between players, the greater the rent that leaders can extract.

The argument that the choice of a solution is a self enforcing contract is quite persuasive but this still leaves unanswered what is the mechanism by which the solution is selected. We take up this in the following two chapters on negotiations.