		Fisherman B	
	no. of boats	1	2
Fisherman A	1	8, 8	1, 6
	2	6, 1	2, 2

FIG. 5.11. A fishing assurance game

to being a 'sucker' (which causes him to receive a payoff of only 1 unit). In short, universal co-operation is the most preferred outcome. Then comes generalized free-riding. Least preferred are those outcomes in which a mismatch of actions occurs. This payoff structure actually determines three possible equilibria, two in pure strategies—each fisherman puts out one boat or each fisherman puts out two boats—and one in mixed strategy. The Pareto-optimal outcome (each fisherman puts one boat out to sea), is only one of the two equilibria in pure strategies. Which equilibrium will be selected actually depends on prior expectations regarding the other's intended action.

Clearly, therefore, the best policy for each party depends on what he thinks the other will do. In actual fact, optimal choice for each fisherman is to put out only one boat if the probability that the other fisherman will choose the same strategy is assessed by him to be in excess of 1/3, and his optimal choice is to put out two boats if this probability is less than 1/3. Denoting by p the probability that the other fisherman puts out one boat, the value of 1/3 is obtained by solving the following equation:

$$8p + 1(1 - p) = 6p + 2(1 - p)$$

which establishes the condition for each fisherman to be indifferent between putting out one boat and putting out two boats to sea. (As is implicit from the above equation, the equilibrium in mixed strategy is such that each fisherman puts out one boat with probability 1/3 and two boats with probability 2/3.)

Thus, there is no certainty that the game will equilibrate at the more favourable of the three (Nash) equilibrium points. It is noteworthy, however, that players need not have complete assurance that others will also co-operate to adopt the same strategy: probabilities significantly smaller than 1 may provide sufficient incentive for co-operation. Still, the possibility exists that the worst equilibrium outcome will emerge *even though* the assumption of common knowledge implies that *each player knows that the other also prefers the co-operative outcome.*[2] This is because there is a genuine trust

[2] Curiously, Taylor rules out this possibility on the grounds that, since both players prefer the co-operative outcome to the mutual defection outcome, 'neither will expect the latter to be the outcome, so the unique Pareto-optimal outcome will result'. The assurance game is consequently deemed to be 'unproblematic' (Taylor, 1987: 39–40). For the reason explained in the text, Taylor's argument is unacceptable. In effect, it comes down to denying the fundamental fact that the 2 × 2 AG comprises three (Nash) equilibria.

problem, that is, a problem of assurance regarding the other person's intended action. Thus, A may know that B would prefer joint co-operation, yet he entertains the fear that B, even though he has corresponding knowledge about his own preference, will choose the maximin strategy ('defect') due to mistrust in what he will himself eventually decide to do. And B can reason in the same way with respect to A's presumed behaviour. *The trust problem is clearly reciprocal since it is basically a problem of mutual expectations*: A may fear that B will abstain from co-operating not because B prefers to free-ride but because B's expectations about his own (A's) behaviour may be pessimistic, and vice versa for B *vis-à-vis* A.

Now, if some form of rudimentary co-ordination device such as pre-play communication (say, in the form of 'cheap talk') is allowed and if the signals sent by the players are interpretable in an unambiguous way, co-operation or joint contribution by both players is much more likely to arise because the players then have the opportunity to reassure one another and to form optimistic expectations about their mutual behaviours.[3] What is worth emphasizing is that the nature of interactions in small groups is highly conducive to pre-play communication and, therefore, if both players' profile is that of an AG player, the Pareto-superior outcome is very likely to be established even in this one-shot game. (Remember that we have reached the same conclusion, *but for repeated games*, when we analysed situations structured as PD.)

Leadership in co-ordination problems

The uncertainty surrounding the players' decisions in a co-ordination problem is overcome as soon as either of the two players can take the initiative in the game with a view to signalling to the other his intention to co-operate. In game-theoretical terms, a particular way of representing the possibility of leadership is by specifying a two-stage assurance game. When the game is played in such a fashion, co-operation by both players is certain to occur: indeed, knowing that the other player will follow suit, the leader has an incentive to make a co-operative move. In other words, by co-operating in the first stage of the game, the leader does not incur the least risk of being 'exploited' by the follower. The outcome (co-operate, co-operate) is clearly a subgame-perfect equilibrium. This is illustrated in Figure 5.12 in which the same payoffs as those assumed in Figure 5.11 have been represented in an extensive form.

If a pure problem of distrust (such as is implicit in the assurance game) can be easily surmounted as soon as one of the players can send a signal or make a first move to the effect that he is determined to co-operate, then, *a fortiori*, the same problem is solved when the game can be repeated. Assume, for instance, that one of the players follows a cautious strategy (start by defecting and, thereafter, co-operating only if the other player has co-operated in the previous round). The other player's best reply to that strategy is obviously not to replicate it but, instead, to start by co-operating (say,

[3] Note that this is precisely the crucial role which Runge ascribes to institutions: to co-ordinate individuals' expectations so as to enable them to co-operate (Runge, 1981, 1984*a*, 1984*b*, 1986).

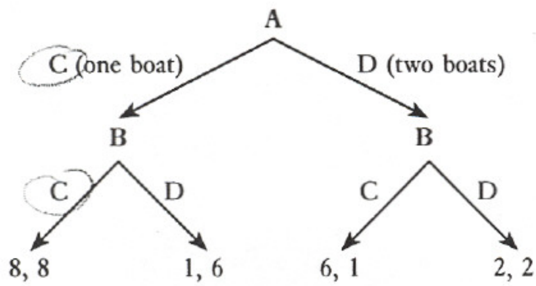

FIG. 5.12. A sequential assurance game

because he follows a strategy of unconditional co-operation) so as to trigger an uninterrupted chain of universal co-operation. Clearly, the cautious strategy is not a Nash equilibrium strategy. However, a 'bad' strategy such as one of unconditional defection is a best reply to itself and therefore supports a Nash equilibrium. (This obviously follows from the fact that, if AG players like best to co-operate, they do not want to be 'exploited'.) What needs to be stressed is that such a strategy is not subgame-perfect since, if by mistake a player co-operates, the other player's best response to that mistake is to co-operate, thus deviating from his Nash equilibrium path. To put it in another way, the commitment of one player to unconditional defection is not credible. Notice that the possibility of a co-operative outcome in such a repeated game is an application of the aforementioned folk theorem and its extension by Benoit and Krishna (1985). Just to give a simple example, a repeated assurance game underlies the observation that in lobster fisheries molesting another fisherman's trap is rarely done because by refraining from doing so a fisherman improves the chances that his own traps will not be molested (Sutinen, Rieser, and Gauvin, 1990: 341).

Threshold effects and freeriding in N*-player co-ordination games*

An interesting feature that arises in connection with co-ordination problems is the existence of threshold effects. As a matter of fact, in many cases, a collective action can bear fruit only if the number of contributors reaches a critical size. To analyse such kinds of situation, we need to consider an N-player assurance game. Let us assume that a given public good (say, the maintenance and management of an irrigation system) yields individual benefits to each member of a group equal to $b(m)$, where m stands for the number of voluntary contributors. Each contributor incurs a fixed cost of c units and, therefore, the total cost for the group is equal to $c \times m$. The choice facing player i can then be represented as in Figure 5.13.

First assume that both $(\partial b(m))/\partial m$ and $(\partial^2 b(m))/\partial m^2$ are positive, implying increasing returns to provision of the public good. Assume also that $b(1) - c < 0$, so that if no other player contributes to the public good, player i also chooses not to contribute. Yet, there exists a critical size m^* such that $b(m) - c > b(m^* - 1)$ or $c < b(m^*) - b(m^* - 1)$: once a certain number, m^*, of other players agree to contribute, player i

Pay-off to player i if the number of other players contributing is

		$n-1$	$n-2$	$n-3$	. . .	0
Player i	contributes	$b(n)-c$	$b(n-1)-c$	$b(n-2)-c$		$b(1)-c$
	does not contribute	$b(n-1)$	$b(n-2)$	$b(n-3)$		0

FIG. 5.13. A N-player assurance game

has an incentive to follow suit since the cost of individual contribution is less than the marginal individual benefit of that contribution. It is evident that, since $(\partial^2 b(m))/\partial m^2 > 0$, if $b(m^*) - c > b(m^* - 1)$, then $b(j) > b(j-1) + c$, $\forall j > m^*$. Therefore, as long as at least m^* other players contribute, player i prefers to co-operate rather than free ride.

In the above game, there are two Nash equilibria in pure strategies. The first equilibrium is characterized by universal defection: given that no one else contributes, player i has no incentive to undertake the collective action alone (we are therefore not in a chicken game). The second equilibrium is characterized by the fact that the collectively optimal level of the public good is provided: everybody contributes to that equilibrium. To avoid falling into the 'bad' equilibrium, a subgroup of players may decide to undertake the collective action in concert, regardless of what the others do. Here lies an important rationale for leadership and the function of the leader consists of mobilizing a sufficient number of contributors rather than showing the good example as assumed in the previous subsection.

It deserves to be noted that an interesting problem which can be raised within the framework of an N-player assurance game is actually a limit case of that analysed above, namely the case in which $b(j) = 0$, $\forall j < n$ and $b(n) > c$. In other words, the collective action can succeed or the public good can be provided only if everyone participates; if only a single agent defects, the public good disappears. The protection of an endangered species or of a breeding-ground illustrates such a possibility that perfectly fits with the description of what an assurance game is about.

Let us now consider the case where there are decreasing returns to scale in the provision of the public good: $(\partial^2 b(m))/\partial m^2$ is negative. In this case, there again exists a critical number of contributors, m^*, below which no individual player has any incentive to contribute. Yet, there now also exists an upper threshold number of contributors, say m^{**}, beyond which the individual marginal benefit of contributing falls short of cost c. The two Nash equilibria in pure strategies are easy to identify: the 'bad' equilibrium in which nobody contributes and a 'nice' equilibrium in which just m^{**} players contribute while the others defect. As long as the size of the group, n, is small (below m^{**}), everyone participates in the collective action under the 'nice' equilibrium. However, in large groups whose size exceeds the threshold m^{**}, the public good is only *partially* produced by a subgroup of players and the amount provided is not Pareto-optimal. It is actually less than the collectively rational amount which would require m^o contributors, with $m^o = \text{argmax}(nb(m^o) - m^o c)$. The collectively rational (co-operative) outcome requires that the collective marginal benefit is equal to the marginal cost c, that is $n(\partial b(m^o))/\partial m^o = c$. It is to be compared to the individually rational (Nash) outcome, m^{**}, which is by definition such that $(\partial b(m^{**}))/\partial m^{**} = c$. Bearing in mind the assumption of decreasing returns to public-good provision, it is evident that $m^o > m^{**}$.

In the latter circumstances (the group is large and $n > m^{**}$), a fraction of the players does not contribute in equilibrium and freeride on the others' efforts. Of course, the wider the gap between the size of the group, n, and the equilibrium threshold number of contributors, m^{**}, the larger the proportion of freeriders. In actual fact, the problem facing the players resembles that of an N-player chicken game, in which the Nash equilibrium would be suboptimal.

In community settings, a large proportion of such freeriders may cause serious tensions to arise. The community may possibly overcome these tensions, however. Thus, it may resort to a co-ordinated solution which has the effect of rotating over time the burden of contributions among the various agents. One option here is to use a correlated equilibrium solution in which contributors are selected through a lottery mechanism. It may also, at a given point of time, ensure that contributors with respect to a given collective action are allowed to abstain from participating in other collective actions so as to distribute equally the costs of public-good provision over a series of different activities. If the above kind of solutions are not applied, an exclusionary process is likely to ensue. This is apparently the case referred to by Ostrom and Gardner (1993) when analysing the Thambesi irrigation system in Nepal. Here, as pointed out earlier, maintenance of the headworks can be carried out by a limited number of the water users and, in particular, the work can be done by head-enders alone. The implication of this situation is that tail-enders may find themselves in a low bargaining position whenever important matters are to be discussed (Ostrom and Gardner, 1993: 97–9).

5.3 Heterogeneous Situations with PD, AG, and CG Players

In real-world settings, groups may not be homogeneous as we have assumed so far. This certainly applies to communities with respect to the management of local-level natural resources. It is indeed often observed that members of a particular user group behave differently because they do not derive the same benefits from a given action. This may be due to a variety of reasons, including differential endowments, different characteristics in terms of the technique employed and the pattern of use of the resource concerned (think of nomadic herders and sedentary agriculturalists), different social identities, different exit possibilities, varying perceptions of the stake involved in resource preservation, etc.

In game-theoretical terms, we will say that, in this case, encounters are heterogeneous in the sense that different *types* of players have to deal with each other. The type of a player is characterized by a particular payoff vector, which may be known or not by the other players. In the following, attention will be focused on heterogeneous games in which players with a payoff structure characteristic of the assurance game face players with a payoff structure characteristic of the prisoner's dilemma. These games are especially interesting because they portray a situation that has much relevance in many human encounters, namely that in which people who do not like to 'exploit' others meet with opportunists. The question that arises in such games is theoretically rich, in so far as it is not a priori clear who among the 'fair' players and the opportunists will determine the final outcome. Before turning to these games, however, mention will be made of two other kinds of heterogeneous encounters. First, we will consider a game in which the two players have an AG payoff structure, yet the benefits accruing to them are not identical. Second, a game in which a player with a chicken game (CG) structure encounters a player with an AG structure will be analysed.

Encounters between two different AG players

Let us assume that the two players who meet in an one-shot game have an AG payoff structure, implying that both of them have no incentive to free–ride on the other's efforts. However, player A has a greater interest in joint co-operation than player B, as illustrated in Figure 5.14.

As usual, there are three Nash equilibria: (C,C), (D,D), and the mixed strategy in which the probability that A plays C is equal to 1/2 and the probability that B plays C is 1/4. Which of these equilibria will emerge depends on the expectations that the players hold about the likelihood that the other co-operates. Assuming that they both hold the same expectation, p, both players co-operate if $p > 1/2$ and defect if $p < 1/4$. Clearly, there exists a range, $1/4 < p < 1/2$, in which A co-operates while B abstains from doing so. Such an outcome, however, is not an equilibrium (player A will not accept to be 'exploited' by player B). If it may arise, it is actually because there exists an inverse relationship between the size of the payoff accruing to the player in case of joint co-operation and the degree of trust required to prompt the player to co-operate.

		Player B	
		C	D
Player A	C	8, 6	1, 5
	D	5, 1	2, 2

FIG. 5.14. A 2 × 2 asymmetrical assurance game

Armenia

$p \cdot 10 + (1-p) 8 = p' (4) + (1-p') 0$

children CG player		AG player C	AG player D
	C	10, 4	8, 4
	D	12, 0	0, 2

FIG. 5.15. A CG player meets an AG player

CG first → (C, C)
AG first → (C, D)

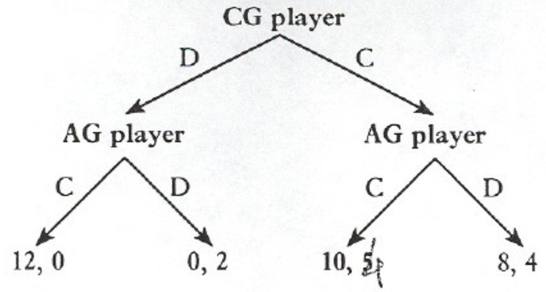

FIG. 5.16. A CG player meets an AG player and leads the sequential game

Encounters between AG and CG players

An interesting situation emerges when an AG player faces a CG player, since it allows us to realize the importance of leadership in determining the equilibrium outcome. To start with, consider the one-shot game with simultaneous moves represented in Figure 5.15.

It is easy to check that, in this game, no equilibrium in pure strategy exists. There is only one equilibrium in mixed strategy, where the probability of the CG player playing C is 2/3 and the probability of the AG player playing C is 4/5.

Equilibria in pure strategies are nevertheless possible as soon as the game is played sequentially. Yet, which equilibrium will arise depends on which player is in the first-mover position. Assume first that the CG player is the leader, as in the sequential game described in Figure 5.16.

It is immediately apparent from Figure 5.16 that joint co-operation will occur: it is in the interest of the CG player to start by co-operating so as to induce the AC player to follow suit. Indeed, if the CG player makes a non-co-operative first move, he is sure to bring about a situation of mutual defection, which he wants absolutely to avoid. In other words, when a party with a leadership role is keen that a collective action is undertaken, but preferably not by himself, whereas the other party tends to follow the leader's behaviour, but prefers bilateral co-operation to bilateral free–riding, joint co-operation will be established. This happy outcome entirely depends on the fact that the leader has a CG payoff structure. Indeed, had the leadership roles been inverted

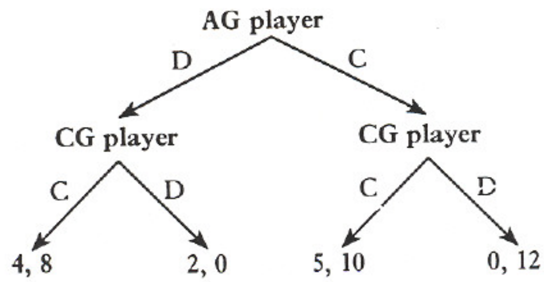

FIG. 5.17. An AG player meets a CG player and leads the game

(the leader is the AG player and the follower is the CG player), universal co-operation would be prevented from arising. This is a straightforward conclusion from Figure 5.17 where the game is led by the AG player.

As is evident from the figure, the collective action will also be undertaken but only by one of the players: having the right to the initial move, the AG player uses the advantage of knowing that the other player has a CG payoff structure to force him to bear the whole cost of this action. Notice carefully, however, that the AG player as a leader is unable to bring about the outcome which he best prefers (universal co-operation) since by co-operating he would incite the follower to defect. This frustration would not have occurred if both players had a CG payoff structure: forcing the other player to co-operate by defecting in the first place is then the ideal outcome which each player wishes for.

An interesting feature which emerges from any encounter between an AG player and a CG player is that both players have an interest in granting leadership to the latter: indeed, the outcome of the first game (Figure 5.16) dominates the outcome of the second game (Figure 5.17). This means that, in a more complex game in which the players would be invited to select the leader before deciding sequentially whether to co-operate or not, the unique subgame-perfect equilibrium path is as follows: the players select the CG player as the leader, thereafter this player co-operates and, in the final stage, the AG player responds by co-operating too. The lesson from such a three-stage game is that, by binding himself to the leadership position, the CG player commits himself to co-operation.

An example which illustrates the aforedescribed situation can again be borrowed from studies of irrigation management. Consider once more a situation in which water users are divided into two subgroups according to whether they are head-enders or tail-enders. Head-enders have a CG payoff structure since they are keen that maintenance of the water control infrastructure is undertaken, but would very much prefer that tail-enders do the work alone (something which may be technically possible, as we have pointed out in the case of the Thambesi irrigation system). On the other hand, tail-enders who are at a locational disadvantage entertain the fear that they may be excluded from decision processes that affect the flow of water reaching their fields (see above): this is why they are eager to participate in maintenance works alongside head-enders, yet would not like to be 'suckers' if head-enders refrain from such partici-

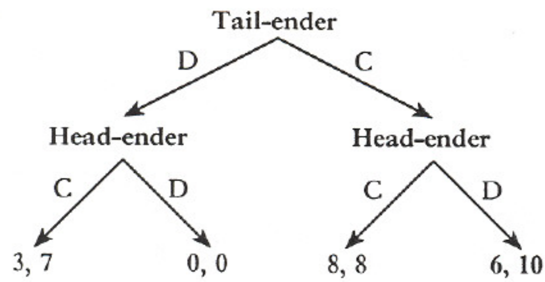

FIG. 5.18. Head-enders deal with tail-enders in a CG-AG encounter

pation (they have an AG payoff structure). In these circumstances, as argued above, head-enders have an incentive to take the leadership so as to associate tail-enders in the maintenance works. This situation seems to fit rather well with the experience of the Pithuwa irrigation project reviewed by Ostrom and Gardner (1993: 105–6).

Another, more plausible situation arises when tail-enders have a mixed CG–AG payoff structure in the following sense: if head-enders participate in the maintenance works, they want to join in order to avoid the aforementioned negative spillover effect but, if head-enders abstain from such participation, they prefer to undertake these works alone rather than leaving the system to fall into decay. If, moreover, as assumed in Figure 5.18, the tail-enders' payoff is higher when they are 'suckers' (6 units) than when they free–ride on the head-enders' maintenance efforts (3 units)—because the price to be paid in terms of loss of reliable access to water is high when participation in these efforts is shunned—the following result obtains: whether the game is played simultaneously or sequentially and whether, in the last case, leadership is exercised by tail-enders or head-enders, the equilibrium outcome is characterized by participation on the part of tail-enders and defection on the part of head-enders. In other words, even though they have the first move, the tail-enders are unable to take advantage of the chicken game structure of the head-ender' payoffs to make them co-operate, owing to the latter's critical control over the supply of water.

Note that, if the tail-ender's payoff when they are 'suckers' is 2 units instead of 6 units or, if their payoff while freeriding on the head-enders' efforts is 6.5 instead of 3 units, it is easy to see that what has just been said is no longer true: if they have the first move, they are now in a position to force head-enders to undertake (alone) the maintenance works. In other words, if the cost of free–riding on maintenance efforts in terms of loss of reliable access to water is not too high for the tail-enders and if they hold a leadership position (two rather implausible assumptions), their leverage allows them to impose the cost of maintenance on the head-enders.

Encounters between AG and PD players in small groups

Let us now turn to the important set of situations in which AG players encounter PD players. There are numerous relevant cases which are worth considering. In order to

PD player

		C	D
AG player	C	2, 2	−1 ~~1~~, 5
	D	1, −1	0, 0

LSO

↑ dominant strategy as CD

FIG. 5.19. A PD player meets an AG player

help the reader to follow the arguments, these cases are presented in increasing order of complexity.

1. To start with, let us examine the simple one-shot two-player game with perfect information in which an AG player meets a PD player. That in this case co-operation cannot occur is immediately evident from the payoff matrix depicted in Figure 5.19.

The non-co-operative outcome (D,D) is the only Nash equilibrium in this game. Note that defection is a dominant strategy for the PD player as a result of which there is no equilibrium in mixed strategy.

2. What would happen if such a game were to be played in a sequential manner? The answer to that question is rather straightforward. If the AG player is in the first-mover position, he will be prompted to defect since he anticipates that the PD player defects in any event. In contrast, and rather unexpectedly, the reverse outcome obtains in the case where the PD player takes the lead: as a matter of fact, knowing that his opponent responds to co-operation by co-operation and to defection by defection, the PD leader has an incentive to start by co-operating. Because (D,D) is the subgame-perfect equilibrium of the former sequential game and (C,C) that of the latter, were both players allowed to choose their leader, both of them would concur in selecting the PD player. Note the similarity between this conclusion and that reached when analysing CG–AG player encounters.

3. If the game is (finitely) repeated rather than being played sequentially, mutual defection appears as the only possible equilibrium as observed in the finitely repeated PD game. Reasoning by backwards induction makes this result clear. As a matter of fact, since the AG player knows that his opponent has a PD payoff structure, he can infer that the latter will surely defect in the last round of the game. His best reply is therefore also to defect in this last round. Being aware that the AG player is going to defect, the PD player has no incentive to build a reputation of 'co-operator' in the round $(t-1)$ and, as a consequence, he will also defect in that round. Since he knows that things will turn out that way, the AG player defects in the same round too. This reasoning can be pursued backwards till the very first round of the game.

4. If the game is played in infinite time (or if its length is finite but indeterminate), the folk theorem applies and joint co-operation is a possible (subgame-perfect) equilibrium.

AG against PD player

	C	D
C	2, 2	−1, 3
D	1, −1	0, 0

AG against AG player

	C	D
C	2, 2	−1, 1
D	1, −1	0, 0

FIG. 5.20. A 2 × 2 game with one-sided asymmetric information

5. By relaxing the assumption of perfect information, we can raise the question as to whether co-operation becomes a possible equilibrium outcome in a finitely repeated game. This question is worth raising bearing in mind the result achieved by Kreps and his associates in the framework of a finitely repeated PD game (see above). Let us first consider the case where uncertainty about the payoff structure of the other player is one-sided asymmetric. More precisely, we assume the first player has a PD payoff structure, and this is common knowledge, while the type of the second player is not known with certainty: the first player assigns probability p to the possibility of the second player being an AG player and probability $(l - p)$ to the possibility that he is a PD player. We know from Kreps *et al.* (1982: 251) that in these circumstances defection will occur throughout the whole game: indeed, if the other player is a PD player, we know that defection is the only possible equilibrium and we have just shown that, if a PD player meets repeatedly an AG player, defection also occurs. There is therefore no reason to expect that a co-operative equilibrium can be generated when there is a doubt about whether the second player has an AG or a PD payoff structure.

6. The next case to consider is that in which the first player is an AG player, and this is common knowledge, while there is doubt about the payoff structure of the second player. Is co-operation more likely to emerge in those more favourable conditions? The answer is a conditional 'yes'. More precisely, co-operation by both players till the last stage of the game may occur if the expectation held by the player with a certain AG payoff structure that the opponent is also an AG player exceeds a certain level. In case this expectation falls below that level, universal defection occurs throughout the game. To see this, consider a three-period game in standard form in which the Row player is known with certainty to be an AG player while there is doubt about whether the Column player is a PD or an AG player. The payoffs pertaining to the two possible kinds of encounters are given in Figure 5.20.

Thus, for instance, by defecting while his opponent co-operates, the PD player gets a payoff of 3 which is more than his payoff when both co-operate. By contrast, in the same circumstances, the AG player gets a payoff of only 1 which is less than his payoff when both co-operate. Let us now assume that an AG player follows the brave reciprocity strategy consisting of starting by co-operating and, thereafter, co-operating as long as the opponent is in 'good standing'. As for the PD player, he starts by co-operating, thereafter mimics what the opponent has done in the previous round till the last round where he defects. The question then is whether these two strategies are the best replies to one another, which would imply that mutual co-operation starts from the first round and continues till the last round when the Column player defects while the Row player co-operates till the very end of the game.

To begin with, consider which payoffs the Column player would earn, *if he is of the AG type*, by following various possible strategies in his encounters with the Row player who is known to be of the AG type and follows the aforedescribed strategy. If he plays a strategy of brave reciprocity (note that similar strategies such as tit for tat or unconditional co-operation also lead to the (C,C,C) sequence of actions), he gets a total payoff over the three periods equal to 6 (2 + 2 + 2). If he plays a strategy whereby, against the brave reciprocity strategy of Row, he co-operates in the first two rounds and defects in the last one, his payoff amounts to 5 (2 + 2 + 1). If he co-operates in the first round and defects in the last two rounds, he earns 3 (2 + 1 + 0) while, if he defects from beginning to end, he earns only 1 (1 + 0 + 0). It is therefore evident that unconditional co-operation dominates the other three strategies. In other words, if Column is of the AG type and Row is of the same type and follows a strategy which consists of starting by co-operating and, thereafter, mimicking what the opponent has done in the previous round, then Column's best reply to the latter is to co-operate throughout the whole game. Since the Row player has adopted the above strategy, he will also co-operate from the beginning to the end of the game. Mutual co-operation therefore occurs till the game ends. There is actually nothing surprising in this result which has already been accounted for at an earlier stage of our analysis.

If Column is of the PD type, his total payoffs while playing various possible strategies against Row are as follows: 6 (2 + 2 + 2) if he plays a co-operative strategy leading to the (C,C,C) sequence of actions; 7 (2 + 2 + 3) if he plays the *fake strategy* whereby he co-operates as his opponent co-operates, and defects in the last round, then revealing his true type; 5 (2 + 3 + 0) if his strategy leads to a (C,D,D) sequence of moves; and 3 (3 + 0 + 0) if he plays unconditional defection. The fake strategy dominates the other strategies available to Column.

Let us now turn to the Row player in order to check whether the brave reciprocity strategy can be a best reply to brave reciprocity (or similar strategies) played by Column if of the AG type and to the fake strategy if of the PD type. Let p be the probability Row assigns to the possibility that Column is of the AG type and $(1-p)$ the probability that Column is of the PD type. If Row plays brave reciprocity, his payoff is 2 for the first period, 2 again for the second period, and $(p \times 2 + (1-p)(-1))$ for the third period, amounting to a total payoff of $3p + 3$ for the three periods together. If, instead, he plays a *safe strategy* whereby he follows the strategy of brave reciprocity except in the last round when he ensures himself against being a 'sucker' by defecting, his payoff is $2 + 2 + (p \times 1 + (1-p)0) = p + 4$. If he plays a strategy leading to the (C,D,D) sequence of actions, he earns $2 + 1 + 0 = 3$ while, if he plays unconditional defection, he gets only 1 $(1 + 0 + 0)$. Clearly, the latter two strategies are dominated by the first two. Whether the first strategy dominates the second strategy or is dominated by it hinges upon the value of p, that is, upon the expectation of Row regarding the